

Randomized Assortment Optimization

Zhengchao Wang, Heikki Peura and Wolfram Wiesemann
Imperial College Business School, Imperial College London, London, UK

When a firm selects an assortment of products to offer to customers, it uses a choice model to anticipate their probability of purchasing each product. In practice, the estimation of these models is subject to statistical errors, which may lead to significantly suboptimal assortment decisions. Recent work has addressed this issue using robust optimization, where the true parameter values are assumed unknown and the firm chooses an assortment that maximizes its worst-case expected revenues over an uncertainty set of likely parameter values, thus mitigating estimation errors. In this paper, we introduce the concept of *randomization* into the robust assortment optimization literature. We show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the firm can improve its worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution. We demonstrate this potential benefit of randomization both theoretically in an abstract problem formulation as well as empirically across three popular choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. We show how an optimal randomization strategy can be determined exactly and heuristically. Besides the superior in-sample performance of randomized assortments, we demonstrate improved out-of-sample performance in a data-driven setting that combines estimation with optimization. Our results suggest that more general versions of the assortment optimization problem—incorporating business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.

1. Introduction

Selecting an assortment of products to offer to customers is a central problem across business operations, with many applications in the travel, hospitality, and retail industries. A firm solving this *assortment optimization problem* seeks to maximize expected revenues or a related objective by selecting a subset of possible products to carry, often subject to constraints such as shelf or display space. With the growth of novel online marketing and e-commerce settings where firms can quickly and flexibly adjust displayed assortments, this classic operations problem is attracting increasing research interest.

The difficulty in this problem lies in accounting for customer behavior: If a prospective customer's preferred product is not part of the offered assortment, she may either substitute it for another one or leave with no purchase. In order to capture this demand substitution,

the firm must specify and estimate a customer *choice model*, which describes a customer's probability of choosing each product for any assortment the firm could offer. Since data on customer behavior in the face of numerous possible assortments is invariably sparse, the estimation of such models is difficult, and estimation errors can lead to significantly suboptimal assortment decisions.

The specification of the choice model involves the familiar bias-variance tradeoff. Most common choice models, such as the multinomial logit (MNL) model (see, *e.g.*, Talluri and van Ryzin 2004), capture customer behavior by imposing strict parametric assumptions. Although this approach results in tractable model estimation and assortment optimization problems that require relatively small amounts of data, imposing restrictive assumptions can introduce a significant bias into estimates (*i.e.*, under-fitting) as well as raise theoretical concerns (such as the independence from irrelevant alternatives, or IIA, property). These pitfalls can be avoided through richer models, such as parametric generalizations of the MNL model, the Markov chain (MC) model (Blanchet et al. 2016), or non-parametric models based on customer preference rankings (Honhon et al. 2012, Bertsimas and Mišić 2019). However, as a result of their added flexibility, more complex models are in turn prone to variance (*i.e.*, over-fitting) unless large amounts of data are available.

The risk of over-fitting the choice model is particularly pernicious when estimates are used to select the best assortment to offer, due to the well-known error-maximization effect of optimization (Smith and Winkler 2006). To address this issue, recent assortment optimization papers have embraced the *robust optimization* paradigm (Rusmevichientong and Topaloglu 2012, Bertsimas and Mišić 2017, Désir et al. 2023). Robust optimization acknowledges the fact that the parameters of the choice model are not known exactly, which is particularly important for complex models with many parameters to be estimated. Instead of point estimates, a decision maker adopting the robust approach specifies an uncertainty set that contains the unknown true parameter values or preferences with a pre-specified confidence. She then chooses an assortment considering the worst parameter setting within this uncertainty set, thus hedging against over-fitting in model estimation.

In this paper, we study *randomized* robust assortment optimization. That is, we show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the decision

maker can improve her worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution.

We first demonstrate the potential benefit of randomization in a general robust assortment optimization problem whose choice model can represent any (reasonable or unreasonable) customer behavior. We derive conditions that the problem must satisfy such that the decision maker benefits from randomizing her choice, that is, such that the problem becomes *randomization-receptive*. We then leverage these conditions to show the benefit of randomization empirically across three popular choice models: the MNL model, the MC model, and the preference ranking (PR) model.

For the ubiquitous MNL model, we show that its randomization-receptiveness hinges on whether the assortment choice is constrained. In the absence of constraints, the decision maker never benefits from randomizing her choice, that is, the model is *randomization-proof*.

By contrast, when the problem is subjected to a cardinality constraint, the MNL model becomes randomization-receptive. We show that not only can the decision maker benefit from choosing an assortment randomly, but the gain from using a randomized strategy over a deterministic one can be arbitrarily large.

The MC model can be randomization-receptive even in the absence of cardinality constraints, but this depends on the characterization of the uncertainty faced by the firm. We show that the MC model is randomization-proof for uncertainty sets that exhibit a specific rectangularity property, which we call product-wise substitution sets. We establish this result through a novel interpretation of the associated assortment optimization problem as a robust Markov decision process (MDP), which allows us to directly apply results from the robust MDP literature. By contrast, for general uncertainty sets the MC model becomes randomization-receptive. Here, the benefits of randomization are bounded in the unconstrained case, but can again become unbounded under a cardinality constraint.

The PR model is also randomization-receptive even in the absence of cardinality constraints. Here again, the benefits of randomization can be unbounded in the constrained setting, whereas they remain bounded in the unconstrained case. We also establish the existence of parsimonious randomization strategies in the PR model: Even though there are exponentially many possible assortments, there are always optimal strategies that randomize between at most $K + 1$ assortments, where K is the number of considered

preference rankings. This parsimony is absent in the constrained MNL and MC models, where all optimal randomization strategies can become arbitrarily complex.

For all three choice models, we also show how an optimal randomization can be determined exactly (by a column generation scheme) and heuristically (through a local search algorithm). We illustrate the runtimes of these solution schemes, as well as the potential benefits of randomization, on synthetic instances of the cardinality-constrained MNL, MC and PR models. For the MNL model, we also demonstrate how the superior in-sample (worst-case) performance of randomized assortments can translate into an improved out-of-sample (expected) performance in a data-driven setting that combines estimation with optimization. Our synthetic experiments are complemented by a data-driven study on a previously published real-life data set, where we show that randomized assortments significantly outperform their deterministic counterparts. A key insight that emerges from our analysis is that more general versions of the assortment optimization problem—where the generality can be owed to the presence of business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.

It is important to recognize that randomization may not be practical in all settings, such as for a small brick-and-mortar retailer with a limited number of stores. The natural application for randomized assortment decisions is in e-commerce and other online settings, where different forms of randomization are already being applied (*e.g.*, in A/B testing various aspects of the sales experience). In these contexts, it is easy for a firm to quickly and flexibly vary the assortment displayed to a customer. Even in conventional retail settings, firms are experimenting with novel strategies adopted from online settings such as A/B testing different aspects of the store design.¹ In a similar vein, a randomized assortment choice could be implemented across a chain of otherwise homogeneous retail stores. Our results show how, in such settings, a firm may benefit from choosing an assortment randomly according to a prudently designed probability distribution.

The remainder of the paper proceeds as follows. We review the relevant related literature in Section 2. We next introduce the nominal, deterministic robust and randomized robust assortment optimization problems, as well as our notions of randomization-receptiveness/-proofness, in Section 3, where we also study the randomization-receptiveness of an abstract

¹ See, *e.g.*, <https://www.bain.com/insights/successful-a-b-tests-in-retail-hinge-on-these-design-considerations/>.

robust assortment optimization problem. Sections 4–6 study the benefits of randomization, as well as exact and heuristic schemes to determine randomized assortments, under the MNL, MC and PR models. Section 7 presents numerical results, and Section 8 concludes the paper. All proofs, the algorithms to compute randomized assortments, as well as details of the numerical experiments are relegated to the appendix.

Notation. We refer to the sets of non-negative and strictly positive real numbers by \mathbb{R}_+ and \mathbb{R}_{++} , respectively. We refer to the vector of all ones and the i -th canonical basis vector as \mathbf{e} and \mathbf{e}_i , respectively; in both cases, the context will dictate the dimension of these objects. For a finite index set \mathcal{X} , we let $\Delta(\mathcal{X}) = \{\mathbf{p} \in \mathbb{R}_+^{|\mathcal{X}|} : \sum_{x \in \mathcal{X}} p_x = 1\}$ denote the associated probability simplex. The Hadamard (element-wise) product is denoted by ‘ \circ ’.

2. Related Literature

This paper is related to the extensive literature on assortment optimization and choice models: Talluri and van Ryzin (2006), Kök et al. (2015), and Gallego and Topaloglu (2019) provide comprehensive overviews of this field. This literature has sought to resolve the twin problems of accurately capturing customer demand substitution using discrete choice models and efficiently finding the corresponding optimal assortment. The most popular choice model is the MNL model dating back to the work of Luce (1959) and Plackett (1975). Although the MNL model is liable to under-fitting data and suffers from the IIA property, it remains popular as both its estimation and the resulting assortment optimization problem can be solved efficiently (Talluri and van Ryzin 2004), even under a cardinality constraint on the size of the offered assortment (Rusmevichientong et al. 2010, Davis et al. 2013).

The literature has proposed a number of richer choice models to account for the MNL model’s shortcomings. These include generalizations of the MNL model, such as the nested logit (Williams 1977, Davis et al. 2014) and mixture of MNL models (Rusmevichientong et al. 2014), which however come at the cost of more difficult estimation and optimization. Recently, two more general classes of choice models have been proposed. Blanchet et al. (2016) develop a tractable Markov chain model that approximates a number of parametric models; a similar idea was used in a simulation study in Zhang and Cooper (2005). Feldman and Topaloglu (2017) consider the MC model in network revenue management, Désir et al. (2020) study the constrained assortment optimization problem, and Şimşek and Topaloglu (2018) propose a method to estimate its parameters. The second class of models is based

on preference rankings, early examples of which include [Mahajan and van Ryzin \(2001\)](#) and [Rusmevichientong et al. \(2006\)](#). This approach considers customer preferences through distributions over preference lists, which allows very general preference structures without imposing a parametric model. [Farias et al. \(2013\)](#) and [van Ryzin and Vulcano \(2015, 2017\)](#) study the estimation of preference ranking models. Although the assortment selection problem is intractable for general PR models ([Aouad et al. 2018](#)), special cases ([Honhon et al. 2012](#), [Paul et al. 2018](#), [Aouad et al. 2021](#)) can be solved efficiently. [Bertsimas and Mišić \(2019\)](#) consider the closely related problem of product line design under this model and propose a mixed-integer optimization based solution approach.

While more complex choice models reduce the bias in estimates, their added flexibility conversely tends to make them prone to variance (over-fitting). This concern is particularly acute when the estimate feeds into the assortment optimization problem due to the error-maximization effect of optimization ([Smith and Winkler 2006](#)), which is well known in finance ([Michaud 1989](#), [DeMiguel and Nogales 2009](#)) and machine learning (see, *e.g.*, [Bishop 2006](#), [Hastie et al. 2009](#)). The robust optimization approach ([Ben-Tal et al. 2009](#), [Bertsimas et al. 2011](#)) explicitly recognizes that estimation should not produce a single point estimate for the choice model parameters but rather an uncertainty set in which the parameters lie with a pre-specified confidence. By selecting the optimal assortment considering the worst parameter setting deemed plausible, the robust approach hedges against over-fitting and can thus be seen as a form of regularization ([El Ghaoui and Lebret 1997](#), [Xu et al. 2009](#)). Robust optimization has been applied to a wide array of operational problems in revenue management ([Birbil et al. 2009](#), [Perakis and Roels 2010](#)), portfolio selection ([Goldfarb and Iyengar 2003](#), [Bertsimas and Sim 2004](#)), inventory management ([Bertsimas and Thiele 2006](#)), facility location ([Baron et al. 2011](#)), and appointment scheduling ([Mak et al. 2015](#)).

The robust approach has proved successful in accounting for parameter uncertainty in choice models. The estimation procedure of [Farias et al. \(2013\)](#) is based on obtaining a worst-case revenues estimate among preference distributions. [Rusmevichientong and Topaloglu \(2012\)](#), [Désir et al. \(2023\)](#), and [Bertsimas and Mišić \(2017\)](#) study the robust assortment optimization problem under the MNL, MC, and PR models, respectively. [Rusmevichientong and Topaloglu \(2012\)](#) introduce uncertainty sets over MNL valuations and

show that this robust MNL model preserves the feature of the nominal model that revenue-ordered assortments are optimal. Désir et al. (2023) extend this robust approach to the MC model and develop efficient algorithms to solve it. Bertsimas and Mišić (2017) consider the related problem of robust product line design under a PR model with both parameter and structural uncertainty.

In all of the aforementioned models, the decision maker deterministically chooses a single assortment to offer. We extend these models by allowing the decision maker to instead randomly choose an assortment according to a probability distribution, and showing when this may benefit her. From a mathematical perspective, the potential benefit from randomization arises from applying robust optimization to a discrete optimization problem (Bertsimas et al. 2016, Delage et al. 2019, Delage and Saif 2022). We show, however, that the superior performance of randomized assortments under this worst-case objective can translate into improved results under the original expected value objective if the model parameters are estimated from data, as is typically the case in practice. Through this insight, as well as our theoretical analysis of the randomization receptiveness of an abstract assortment optimization problem, we also contribute to the robust optimization literature.

While randomized strategies appear to be new in the context of assortment optimization, they have been successfully applied in a growing number of related revenue management contexts. The most prominent example of this are heuristic admission control strategies employed in network revenue management (Reiman and Wang 2008, Jasin and Kumar 2012, Jasin 2015, Bumpensanti and Wang 2020) to decide which assortment (*e.g.*, of flight legs) to offer to customers over time (or which classes of customers to admit) under uncertainty and resource constraints. These papers show that judiciously resolving a deterministic approximation of the problem can provide guarantees on expected revenues losses, when the firm probabilistically decides which assortment to offer, which is referred to as the allocation control (PAC) heuristic. Ferreira et al. (2018) study a retailer’s dynamic pricing problem with inventory constraints and learning about customer demand. The retailer balances an exploration-exploitation tradeoff by selecting a distribution over a discrete set of prices to offer. In a recent paper, Ma (2023) compares the decision to offer an assortment with the use of lotteries where participating buyers do not know which product they will eventually receive. Drawing on the mechanism design literature, he provides conditions for offering an assortment to be optimal compared to such randomized allocation mechanisms.

We add to this emerging stream of literature by showing the value of optimal randomized strategies in the classic assortment optimization problem.

3. Value of Randomization in Robust Assortment Optimization

In this section, we define the generic robust assortment optimization problem to study the value of randomized strategies in solving this problem. In Section 3.1 we study *whether* a randomized strategy may benefit a firm solving this generic problem, and derive conditions for this to be the case. Section 3.2 then discusses the intuition for *why* randomization can be beneficial. In the subsequent Sections 4-6, we investigate *when*, under prevalent choice models, randomization offers practical benefits.

Section 3.3 summarizes these results.

In the *nominal assortment optimization problem*, a firm chooses an assortment to offer to customers out of N candidate products. A customer then either buys one of the offered products or makes no purchase, according to a discrete choice model. Formally, an instance of the problem is a tuple $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathbf{r})$ defined as follows. Let $\mathcal{N} = \{1, \dots, n\}$ denote the set of products, and let $\mathcal{N}_0 = \mathcal{N} \cup \{0\}$ be the extended set that contains the no-purchase alternative indexed 0. We denote by $\mathcal{S} \subseteq \{S : S \subseteq \mathcal{N}\}$ the set of admissible assortments, which may exclude some subsets of \mathcal{N} due to (*e.g.*, cardinality) constraints. When the problem is unconstrained, \mathcal{S} is the power set of \mathcal{N} . The choice model is a mapping $\mathfrak{C} : \mathcal{S} \rightarrow \Delta(\mathcal{N}_0)$ such that $\mathfrak{C}(i|S) = 0$ for all $i \in \mathcal{N} \setminus S$, $S \in \mathcal{S}$. The price of product $i \in \mathcal{N}$ is $r_i > 0$; we also include the price of the no-purchase option, $r_0 = 0$, in the vector of product prices $\mathbf{r} \in \mathbb{R}_+^{n+1}$. The firm chooses an assortment $S^* \in \mathcal{S}$ that maximizes its expected revenues $R(S) = \sum_{i \in \mathcal{N}} r_i \cdot \mathfrak{C}(i|S)$ among all admissible assortments $S \in \mathcal{S}$:

$$R_{\text{nom}}^* = \max_{S \in \mathcal{S}} R(S). \quad (\text{NOMINAL})$$

In reality, the choice model is typically estimated from data and hence uncertain. The *robust assortment optimization problem* takes this uncertainty into account. An instance of this problem is a tuple $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \mathbf{r})$, where the set of products \mathcal{N} and the set of admissible assortments \mathcal{S} are defined as before. The choice model $\mathfrak{C} : \mathcal{S} \times \mathcal{U} \rightarrow \Delta(\mathcal{N}_0)$, however, now depends on a realization $u \in \mathcal{U}$ from an uncertainty set \mathcal{U} (*e.g.*, over choice model parameters). To avoid technicalities, we assume that \mathcal{U} is a compact subset of a finite-dimensional space. We stipulate that $\mathfrak{C}(i|S, u) = 0$ for all $i \in \mathcal{N} \setminus S$, $S \in \mathcal{S}$ and $u \in \mathcal{U}$. The *deterministic*

robust assortment optimization problem (see, e.g., Rusmevichientong and Topaloglu 2012) then chooses the assortment that maximizes the worst-case expected revenues

$$R_{\text{det}}^*(\mathcal{U}) = \max_{S \in \mathcal{S}} R^*(S), \quad (\text{DETERMINISTIC ROBUST})$$

where $R^*(S) = \min_{u \in \mathcal{U}} R(S, u)$ and $R(S, u) = \sum_{i \in S} r_i \cdot \mathfrak{C}(i|S, u)$. When the uncertainty set $\mathcal{U} = \{u^0\}$ is a singleton, the choice model is known exactly and **DETERMINISTIC ROBUST** recovers the classical **NOMINAL** problem.

Instead of selecting a single assortment to offer, the firm could offer *each* assortment $S \in \mathcal{S}$ with a probability $p_S \geq 0$. In this paper, we therefore propose to solve the *randomized robust assortment optimization problem*:

$$R_{\text{rand}}^*(\mathcal{U}) = \max_{\mathbf{p} \in \Delta(\mathcal{S})} R^*(\mathbf{p}), \quad (\text{RANDOMIZED ROBUST})$$

where $R^*(\mathbf{p}) = \min_{u \in \mathcal{U}} R(\mathbf{p}, u)$ and $R(\mathbf{p}, u) = \sum_{S \in \mathcal{S}} p_S \cdot R(S, u)$. The deterministic assortments $S \in \mathcal{S}$ correspond to degenerated randomized assortments $\mathbf{e}_S \in \Delta(\mathcal{S})$ that place all probability mass onto individual assortments, and thus the feasible region of **RANDOMIZED ROBUST** contains the feasible solutions of **DETERMINISTIC ROBUST**. Note that depending on its argument, R may refer either to the nominal expected revenues or the expected revenues under the parameter realization $u \in \mathcal{U}$, and that R^* may refer to the worst-case expected revenues of a single assortment or a randomization strategy.

3.1. Can Randomization Be Beneficial?

We now study whether the decision maker in the robust assortment optimization problem can benefit from randomized strategies for generic choice models \mathfrak{C} and uncertainty sets \mathcal{U} . In later sections, we consider the value of randomization under three common choice models: the MNL model, the MC model, and the PR model, each of which implies different definitions for \mathfrak{C} and \mathcal{U} .

Since the **RANDOMIZED ROBUST** problem subsumes the **DETERMINISTIC ROBUST** problem, clearly $R_{\text{rand}}^*(\mathcal{U}) \geq R_{\text{det}}^*(\mathcal{U})$. Our goal is to show under what conditions randomization can benefit the firm, that is, when the strict inequality $R_{\text{rand}}^*(\mathcal{U}) > R_{\text{det}}^*(\mathcal{U})$ is satisfied. We say the problem is then receptive to randomization, as per the following definition.

DEFINITION 1 (RANDOMIZATION-RECEPTIVENESS/PROOFNESS). An instance of the robust assortment optimization problem is *randomization-receptive* if $R_{\text{rand}}^*(\mathcal{U}) > R_{\text{det}}^*(\mathcal{U})$

and *randomization-proof* otherwise. Likewise, we say that the (un-)constrained robust assortment optimization problem under a particular choice model is randomization-proof if all of its instances are randomization-proof, and it is randomization-receptive otherwise.

We note that **NOMINAL** is randomization-proof by construction. Indeed, for singleton uncertainty sets $\mathcal{U} = \{u^0\}$, **RANDOMIZED ROBUST** reduces to a linear program that attains its optimal value at an extreme point of the probability simplex $\Delta(\mathcal{S})$. This in turn corresponds to a randomization that places unit probability onto a single assortment $S \in \mathcal{S}$.

In contrast, the decision maker can benefit from randomization in the robust assortment optimization problem. The following result establishes when this is the case by providing a necessary and sufficient condition for an instance of the problem to be randomization-proof.

THEOREM 1. *Fix a robust assortment optimization instance $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \mathbf{r})$. The instance is randomization-proof if and only if there is $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R_{\text{det}}^*(\mathcal{U})$ for all $S \in \mathcal{S}$, where $\mathcal{P}(\mathcal{U})$ denotes the set of all probability distributions supported on \mathcal{U} .*

We illustrate the statement of **Theorem 1** with an example. Consider a robust assortment optimization problem with three feasible assortments S_1 , S_2 and S_3 as well as two possible uncertainty realizations u and u' . The expected revenues of each assortment are $(R(S_1, u), R(S_2, u), R(S_3, u)) = (2, 0, 1)$ and $(R(S_1, u'), R(S_2, u'), R(S_3, u')) = (1/2, 3/2, 1)$. Figure 1 (left) shows that this problem is *randomization-proof*. The figure depicts the expected revenues for all possible randomizations over the three assortments for u and u' , with the worst-case expected revenues shown as solid regions. The optimal deterministic solution is $\mathbf{e}_{S_3} = (0, 0, 1)$, located along the vertical axis, with $R_{\text{det}}^*(\mathcal{U}) = 1$. Note that any deviation from the assortment S_3 leads to a deterioration of the worst-case expected revenues: there is no direction in which we can move away from \mathbf{e}_{S_3} to improve the worst-case expected revenues. In terms of the condition in **Theorem 1**, the optimality of the deterministic solution \mathbf{e}_{S_3} is certified by the distribution $\mathbb{Q} = \frac{1}{3} \cdot \delta_u + \frac{2}{3} \cdot \delta_{u'}$, so that the expected value $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})]$ equals $R_{\text{det}}^*(\mathcal{U}) = 1$ for all three assortments.

Next consider the same example but with the updated expected revenues $(R(S_1, u), R(S_2, u), R(S_3, u)) = (1, 1/2, 3/2)$ and $(R(S_1, u'), R(S_2, u'), R(S_3, u')) = (1/2, 2, 1)$. Figure 1 (right) shows that this problem is *randomization-receptive*. In contrast to the previous case, a direction to improve the worst-case revenues now exists at each deterministic solution \mathbf{e}_{S_1} , \mathbf{e}_{S_2} and \mathbf{e}_{S_3} . In terms of **Theorem 1**, any distribution \mathbb{Q}

with $\mathbb{Q}[\tilde{u} = u'] < 1$ implies that $\mathbb{E}_{\mathbb{Q}}[R(S_3, \tilde{u})] > R_{\text{det}}^*(\mathcal{U})$, while the distributions satisfying $\mathbb{Q}[\tilde{u} = u'] = 1$ imply $\mathbb{E}_{\mathbb{Q}}[R(S_2, \tilde{u})] = 2 > R_{\text{det}}^*(\mathcal{U})$.

Together, the two graphs in Figure 1 suggest that a deterministic assortment S maximizes the worst-case expected revenues if and only if the graph $\mathbf{p} \mapsto R^*(\mathbf{p})$ does not contain a feasible direction of improvement at the point \mathbf{e}_S . This is indeed what the condition in Theorem 1 amounts to: At every optimal deterministic assortment, there is no direction to move without compromising worst-case expected revenues. This intuition can be made precise by considering the super-differential of $R^*(\mathbf{e}_S)$, and the resulting condition turns out to be equivalent to the condition in Theorem 1. We relegate the formal justification of this statement to Appendix B.




Figure 1 Randomization-proof (left) and randomization-receptive (right) instances of **RANDOMIZED ROBUST**. In both cases, we show two-dimensional projections of the three-dimensional probability simplex $\Delta(S)$ with $S = \{S_1, S_2, S_3\}$ (grey triangles). The red and blue planes illustrate the expected revenues under the two uncertainty realizations u and u' , respectively; solid (hollow) areas indicate regions where the respective realization constitutes the worst (best) case.

The condition in Theorem 1 closely relates to the concept of strong duality, which we will use in subsequent sections when discussing common choice models. Specifically, strong duality implies the condition in Theorem 1, but the reverse is not true. Consequently, as the next result shows, strong duality implies randomization-proofness.

COROLLARY 1. *Suppose that strong duality holds for the problem $(\mathcal{N}, \mathcal{S}, \mathcal{C}, \mathcal{U}, \mathbf{r})$:*

$$\max_{S \in \mathcal{S}} \min_{u \in \mathcal{U}} R(S, u) = \min_{u \in \mathcal{U}} \max_{S \in \mathcal{S}} R(S, u).$$

Then the problem is randomization-proof.

Intuitively, any optimal deterministic assortment achieves the objective value of the left-hand side of the equation in [Corollary 1](#), while the right-hand side of this equation constitutes an upper bound on the worst-case expected revenues achievable by any (deterministic or randomized) assortment strategy: It records the expected revenues that can be achieved if the decision maker knew the uncertainty realization u prior to making an assortment choice. Since this ‘crystal ball’ upper bound is achieved by an optimal deterministic assortment, the decision maker cannot improve by randomizing between assortments. The reverse statement, however, is not true: A problem instance may be randomization-proof even when strong duality does not hold. Indeed, the left-hand example in [Figure 1](#) illustrates this: we can verify that the ‘crystal ball’ solution is u' with expected revenues 1.5 (at S_2), which is greater than $R_{\text{det}}^*(\mathcal{U}) = 1$. Thus, strong duality does not hold, yet our earlier discussion showed that the instance is randomization-proof.

The equivalence between strong duality and the condition in [Theorem 1](#) holds in the special case where the worst-case uncertainty realization is unique. [Theorem 1](#) then simplifies to the following result: The problem is randomization-proof if and only if, under the worst-case realization, the decision maker cannot improve her expected revenues by deviating from the optimal deterministic assortment S^* .

COROLLARY 2. *Fix a robust assortment optimization problem $(\mathcal{N}, \mathcal{S}, \mathfrak{C}, \mathcal{U}, \mathbf{r})$ and let S^* be an optimal deterministic assortment. If the worst-case parameter set $\arg \min_{u \in \mathcal{U}} R(S^*, u)$ is a singleton $\{u^*\}$, then the problem is randomization-proof if and only if*

$$R^*(S^*) \geq \max_{S \in \mathcal{S}} R(S, u^*).$$

3.2. Why Is Randomization Beneficial?

[Theorem 1](#) and [Corollary 2](#) provide conditions under which the firm benefits from randomizing its assortment. We now investigate why this is the case. To establish an intuition, consider the following example.

EXAMPLE 1. *A firm wishes to offer an assortment containing no more than two out of a product universe of three products, with prices $r_1 = 10$, $r_2 = 9$ and $r_3 = 8$. The customers’ purchase behavior follows a multinomial logit model with nominal product valuations $\mathbf{v}^0 = (v_0^0, v_1^0, v_2^0, v_3^0) = (1, 1, 1, 1)$, where the outside option has value $v_0^0 = 1$. The optimal nominal assortment under the product valuations \mathbf{v}^0 comprises the products $\{1, 2\}$ and generates expected revenues of 6.33 (cf. the first row of [Table 1](#)).*

Valuation	Assortments						
	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	$0.476 \cdot \{1,2\} + 0.524 \cdot \{2,3\}$
v^0	5.00	4.50	4.00	6.33	6.00	5.67	5.98
v^1	1.67	5.09	5.33	5.48	5.63	6.44	5.98
v^2	7.5	3.00	3.56	7.67	7.58	4.74	6.13

Table 1 Expected revenues of different deterministic as well as the optimal randomized assortment under the three valuations v^0 , v^1 and v^2 in our motivating example. The green and red shadings refer to the nominal and worst-case performance of the optimal nominal assortment, respectively; the purple and cyan shadings refer to the worst-case performances of the optimal deterministic robust and randomized robust assortments, respectively.

Assume now that the product valuations may deviate from v^0 and instead amount to $v^1 = (v_0^1, v_1^1, v_2^1, v_3^1) = (1, 0.2, 1.3, 2)$ or $v^2 = (v_0^2, v_1^2, v_2^2, v_3^2) = (1, 3, 0.5, 0.8)$ as well. In that case, Table 1 shows that the expected revenues of the optimal nominal assortment $\{1, 2\}$ decrease by up to 13.47% to 5.48 under scenario v^2 . A firm using the deterministic robust approach would instead offer the assortment comprising the products $\{1, 3\}$, generating worst-case expected revenues of 5.63 that are 2.29% higher than those of the optimal nominal assortment. However, by prudently randomizing between the assortments $\{1, 2\}$ and $\{2, 3\}$, the firm can further increase its worst-case expected revenues by 5.67% to 5.98.

The benefits of randomization can be understood from three complementary perspectives. In the example, the firm faces the risk of adverse outcomes due to taking decisions under parameter uncertainty. From a managerial perspective, randomization allows the firm to diversify against this risk and alleviate the error-maximizing effect of optimization. From a mathematical viewpoint, this diversification amounts to enlarging the feasible region of the deterministic robust assortment optimization problem by embedding the discrete assortment choices of the latter into a probability simplex that allows the firm to select convex combinations of multiple assortments in the randomized problem. From a game-theoretic viewpoint, finally, the deterministic robust assortment optimization problem constitutes a Stackelberg leader-follower game where the firm (the leader) chooses an assortment and ‘nature’ (the follower) picks the worst valuation vector v^i upon observing the selected assortment. In the randomized problem, on the other hand, nature only observes the firm’s randomization strategy (as opposed to the randomly selected assortment), which restricts its power. As per Theorem 1 and Corollary 2, the firm then enjoys diversification benefits from randomization as long as a profitable move away from the optimal deterministic assortment exists under the worst-case realization. Indeed, the example also illustrates the randomization-proofness condition in Corollary 2: Since each optimal

Table 2 Summary of theoretical results on the benefits of randomization.

Choice model	Constrained	Randomization receptiveness	Maximum benefit $R_{\text{rand}}^*(\mathcal{U})/R_{\text{det}}^*(\mathcal{U})$	Complexity of randomization
Multinomial Logit Section 4	No	No	—	—
	Yes	Yes Theorem 2	Unbounded Theorem 2	Exponential Theorem 2
Markov Chain Section 5	No	Yes/No Theorem 4, Observation 2	Bounded Observation 3	(Unknown)
	Yes	Yes Theorem 5	Unbounded Theorem 5	Exponential Theorem 5
Preference Ranking Section 6	No	Yes Observation 4	Bounded Observation 5	Linear Theorem 7
	Yes	Yes Theorem 6	Unbounded Theorem 6	Linear Theorem 7

deterministic assortment in the table has a unique worst-case valuation, the condition in [Corollary 2](#) does not hold, and the problem is randomization-receptive.

3.3. When Is Randomization Beneficial?

In [Sections 4-6](#), we study the benefits of randomization under three prevalent choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. For each model, we investigate *(i)* whether the problem can benefit from randomization; *(ii)* how large the potential benefits are; and *(iii)* how complex the optimal randomization strategies may be, *i.e.*, how many different assortments they may involve.

[Table 2](#) summarizes our findings. The randomization-receptiveness of the robust assortment optimization problem depends on the complexity of the choice model and underlying business constraints. Under the parsimonious MNL model, the problem is randomization-proof in the absence of any further constraints on the chosen assortment. But when we either impose a cardinality constraint on the size of the assortment, or use the MC or PR models to capture more complex customer behavior, randomization becomes beneficial. Investigating the magnitude of the potential benefits of randomization, we find that these are bounded in the absence of business constraints, but they can become arbitrarily large under cardinality constraints. Interestingly, under both the MNL and MC models, the optimal randomization strategy may involve exponentially many assortments, whereas the PR model always admits parsimonious optimal randomization strategies that involve a small number of assortments.

4. Multinomial Logit Model

We first study the robust assortment optimization problem under the MNL model. We show that whether the firm can benefit from randomization depends on the existence of a cardinality constraint on the size of the assortment. The unconstrained problem is randomization-proof, while the cardinality-constrained problem is not only randomization-receptive, but the resulting benefit can also be arbitrarily large (Section 4.1). Section 4.2 develops algorithms for optimally solving the latter problem.

The MNL model is parameterized by a vector of customer product valuations $\mathbf{v} = (v_0, v_1, \dots, v_n) \in \mathbb{R}_{++}^{n+1}$. Given these valuations and an assortment $S \in \mathcal{S}$, a customer purchases product $i \in \mathcal{N}$ with probability

$$\psi_i(S, \mathbf{v}) = \begin{cases} \frac{v_i}{v_0 + \sum_{j \in S} v_j} & \text{if } i \in S, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

and the corresponding no-purchase probability is $\psi_0(S, \mathbf{v}) = 1 - \sum_{i \in S} \psi_i(S, \mathbf{v})$. The expected revenues amount to

$$R(S, \mathbf{v}) = \sum_{i \in S} r_i \cdot \psi_i(S, \mathbf{v}) = \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i}, \quad (2)$$

where we identify the parameter vector \mathbf{u} with the product valuations \mathbf{v} . In the corresponding robust assortment optimization problems, we assume that \mathbf{u} is only known to be contained in a compact uncertainty set $\mathcal{U} = \mathcal{V} \subseteq \mathbb{R}_{++}^{n+1}$.

4.1. Randomization-Receptiveness of the Multinomial Logit Problem

The unconstrained problem presented above is randomization-proof. Indeed, [Désir et al. \(2023\)](#) show that the problem satisfies strong duality, that is, we have

$$\max_{S \in \mathcal{S}} \min_{\mathbf{v} \in \mathcal{V}} R(S, \mathbf{v}) = \min_{\mathbf{v} \in \mathcal{V}} \max_{S \in \mathcal{S}} R(S, \mathbf{v}).$$

[Corollary 1](#) then implies that randomization over assortments cannot benefit the firm.

Strong duality, however, ceases to hold if we impose a cardinality constraint $|S| \leq C$ on the size of the offered assortment. We have seen in [Example 1](#) that the firm may benefit from randomization under a cardinality constraint. We now show that the benefit from randomization can indeed be arbitrarily large in the constrained problem, but that at the same time the optimal strategy may require randomization between many assortments.

THEOREM 2. *For any number of products $n \geq 2$ and any restriction $|S| \leq C$, $C \in \{1, \dots, n-1\}$, there are instances of the cardinality-constrained robust MNL problem where*

1. $R_{\text{det}}^*(\mathcal{V}) = 0$ while $R_{\text{rand}}^*(\mathcal{V}) > 0$;
2. *the unique optimal randomized assortment strategy places equal (positive) probability on each assortment S satisfying $|S| = C$ and zero probability on all other assortments.*

The first part of [Theorem 2](#) implies that the benefits $R_{\text{rand}}^*(\mathcal{V})/R_{\text{det}}^*(\mathcal{V})$ from randomization are unbounded. The proof of the theorem, which is relegated to the appendix, considers a class of robust constrained MNL instances where all products have equal price and the uncertainty set \mathcal{V} comprises all valuation vectors \mathbf{v} in which exactly B valuations are zero and the remaining valuations are 1. For $B \geq C$, where C is the admissible assortment cardinality, any deterministic assortment S results in worst-case expected revenues of zero since the uncertainty set contains a vector \mathbf{v} in which $v_i = 0$ for all $i \in S$. The optimal randomized strategy, on the other hand, places equal (positive) probability on all assortments with exactly C products, and zero probability on all other assortments. As long as $B < n$, this randomized strategy raises strictly positive worst-case revenues since there is a strictly positive probability that the offered assortment contains products whose values are strictly positive even under the worst-case valuation realization. The optimal randomization strategy, however, may involve exponentially many assortments, as we can see from the second statement by setting $C = \lfloor n/2 \rfloor$, which results in $\Theta(2^n/\sqrt{n})$ many assortments. Optimal randomization strategies can thus generally be complex. We next turn to the problem of finding such optimal strategies.

4.2. Solving the Randomized Constrained Multinomial Logit Problem

[Appendix A](#) presents two algorithms to compute randomization strategies for the constrained robust MNL problem with a binary representable uncertainty set of the form

$$\mathcal{V} = \{\mathbf{v} = \mathbf{F}\boldsymbol{\xi} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}, \boldsymbol{\xi} \in \{0, 1\}^m\}, \quad (3)$$

where $\mathbf{F} \in \mathbb{R}^{(n+1) \times m}$, $\mathbf{A} \in \mathbb{R}^{l \times m}$ and $\mathbf{b} \in \mathbb{R}^l$. While any discrete uncertainty set is binary representable, we list some popular uncertainty sets that enjoy compact representations.

- (i) **Budget uncertainty sets.** For lower and upper valuation bounds $\underline{\mathbf{v}}, \bar{\mathbf{v}} \in \mathbb{R}^{n+1}$ and an uncertainty budget $\Gamma \in \mathbb{N}$, we define the uncertainty set

$$\mathcal{V} = \{\mathbf{v} = \bar{\mathbf{v}} - (\bar{\mathbf{v}} - \underline{\mathbf{v}}) \circ \boldsymbol{\xi} : \mathbf{e}^\top \boldsymbol{\xi} \leq \Gamma, \boldsymbol{\xi} \in \{0, 1\}^{n+1}\}.$$

Under the budget uncertainty set, up to Γ valuations v_i can attain their lower bounds \underline{v}_i , whereas the remaining valuations v_i attain their upper bounds \bar{v}_i .

(ii) **Factor model uncertainty sets.** For a nominal valuation vector $\mathbf{v}^0 \in \mathbb{R}^{n+1}$ and a factor loading matrix $\Phi \in \mathbb{R}^{(n+1) \times m}$, we define the uncertainty set

$$\mathcal{V} = \{ \mathbf{v} = \mathbf{v}^0 + \Phi(2\xi - \mathbf{e}) : \xi \in \{0, 1\}^m \}.$$

Here, the valuations \mathbf{v} differ from their nominal values \mathbf{v}^0 by $\Phi\mathcal{B}_\infty$, where $\mathcal{B}_\infty = \{-1, 1\}^m$ contains the extreme points of the unit ∞ -norm ball in \mathbb{R}^m .

(iii) **Norm uncertainty sets.** For a nominal valuation vector $\mathbf{v}^0 \in \mathbb{R}^{n+1}$ and lower and upper valuation bounds $\underline{\mathbf{v}}, \bar{\mathbf{v}} \in \mathbb{R}^{n+1}$, we define the uncertainty set

$$\mathcal{V} = \left\{ \mathbf{v} = \mathbf{v}^0 + (\bar{\mathbf{v}} - \mathbf{v}^0) \circ \xi^+ + (\underline{\mathbf{v}} - \mathbf{v}^0) \circ \xi^- : \begin{array}{l} \|(\bar{\mathbf{v}} - \mathbf{v}^0) \circ \xi^+ + (\underline{\mathbf{v}} - \mathbf{v}^0) \circ \xi^-\|_p \leq \theta \\ \xi^+ + \xi^- \leq \mathbf{e}, \xi^+, \xi^- \in \{0, 1\}^{n+1} \end{array} \right\}$$

where $p \in \{0, 1, \infty\}$ and $\theta \in \mathbb{R}_+$. Norm uncertainty sets hedge against valuations \mathbf{v} that reside in a θ -neighbourhood of the nominal valuations \mathbf{v}^0 , as measured by the p -norm.

The randomized constrained MNL model can be formulated as a robust linear program. To this end, denote by $\mathcal{S} = \{S \in \mathcal{N} : |S| \leq C\}$ the set of assortments whose cardinality is less or equal to the size restriction C . The problem can then be written as

$$\begin{aligned} & \text{maximize} && \min_{\mathbf{v} \in \mathcal{V}} \sum_{S \in \mathcal{S}} p_S \cdot \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i} \\ & \text{subject to} && \sum_{S \in \mathcal{S}} p_S = 1 \\ & && p_S \geq 0, S \in \mathcal{S}. \end{aligned} \tag{4}$$

This problem is computationally challenging since it typically comprises an exponential number of decision variables, and its objective function contains an embedded optimization problem that minimizes a non-convex function over a discrete uncertainty set \mathcal{V} . For later reference, we note that the dual of problem (4) amounts to the robust linear program

$$\begin{aligned} & \text{minimize} && \max_{S \in \mathcal{S}} \sum_{\mathbf{v} \in \mathcal{V}} \kappa_{\mathbf{v}} \cdot \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i} \\ & \text{subject to} && \sum_{\mathbf{v} \in \mathcal{V}} \kappa_{\mathbf{v}} = 1 \\ & && \kappa_{\mathbf{v}} \geq 0, \mathbf{v} \in \mathcal{V}. \end{aligned} \tag{5}$$

Strong duality between (4) and (5) holds since (4) is feasible by construction. Similar to problem (4), problem (5) is computationally challenging due to the typically exponential number of decision variables as well as the embedded maximization over a discrete ‘uncertainty set’ \mathcal{S} whose decision variables S appear in a non-convex objective function.

In Appendix A, we propose an exact column generation scheme as well as a heuristic local search method to solve the randomized constrained MNL problem.

5. Markov Chain Model

The MC model proposed by Blanchet et al. (2016) generalizes the MNL model. Under the MC model, the choice behavior of a customer is described through a vector $\boldsymbol{\lambda} \in \mathbb{R}_+^{n+1}$, where λ_i , $i = 0, \dots, n$, denotes the probability of product i being the most preferred choice, as well as a matrix $\boldsymbol{\rho} = (\rho_{ij})_{i,j} \in \mathbb{R}_+^{(n+1) \times (n+1)}$, where ρ_{ij} , $i, j = 0, \dots, n$, characterizes the probability of the customer substituting product i by product j if product i is not available. In other words, an arriving customer attempts to purchase each product $i = 0, \dots, n$ with probability λ_i . If the preferred product, say i , is not offered, then she attempts to purchase each product $j = 0, \dots, n$ with probability ρ_{ij} , and the process continues as if j had been her preferred choice. We require that $\rho_{00} = 1$ as well as $\rho_{ii} = 0$ and $\rho_{i0} > 0$ for all $i \neq 0$.

In the robust assortment optimization problem under the MC model (Désir et al. 2023), the substitution matrices $\boldsymbol{\rho}$ are only known to reside in an uncertainty set

$$\mathcal{U} \subseteq \left\{ \boldsymbol{\rho} \in \mathbb{R}_+^{(n+1) \times (n+1)} : \sum_{j=0}^n \rho_{ij} = 1 \quad \forall i = 0, \dots, n \right\},$$

which we assume to be compact in order to avoid technicalities. Under this setting, the substitution behavior of the customers can be determined by any substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$, and the decision maker optimizes the expected revenues considering the worst such matrix. We require that all $\boldsymbol{\rho} \in \mathcal{U}$ satisfy $\rho_{00} = 1$ as well as $\rho_{ii} = 0$ and $\rho_{i0} > 0$ for all $i \neq 0$. In addition, we require that there exists a $\boldsymbol{\rho} \in \mathcal{U}$ and a product $i \neq 0$ satisfying that $\rho_{i0} < 1$.

We now study whether the robust assortment optimization problem under the MC model benefits from randomization. To do so, we first express the problem as a robust Markov decision process (MDP) in Section 5.1. Leveraging this interpretation, we then show that the benefits of randomization depend on the nature of parameter uncertainty. Specifically, we contrast product-wise substitution sets where no information is available about the dependence of ρ_{ij} and ρ_{kl} whenever $i \neq k$ and general substitution sets (Section 5.2). Section 5.3, finally, discusses the solution of the (un-)constrained robust assortment optimization problem under the MC model.

5.1. A Robust MDP Reformulation for the Unconstrained Markov Chain Problem

Our objective is to reformulate the robust assortment optimization problem under the MC model as an instance of a robust MDP as per the following definition.

DEFINITION 2 (ROBUST MDP). A robust MDP is defined by the tuple $(\mathcal{X}, \mathcal{A}, \mathcal{P}, q, c, \gamma)$, where \mathcal{X} denotes the state space, \mathcal{A} represents the action space, $q(x)$, $x \in \mathcal{X}$, characterize the initial state distribution, $c(x, a)$, $x \in \mathcal{X}$ and $a \in \mathcal{A}$, denote the immediate rewards, and $\gamma \in (0, 1)$ is the discount factor. The ambiguity set

$$\mathcal{P} \subseteq \left\{ p : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \rightarrow \mathbb{R}_+ : \sum_{x' \in \mathcal{X}} p(x'|x, a) = 1 \quad \forall x \in \mathcal{X}, \forall a \in \mathcal{A} \right\}$$

contains all transition kernels p that are deemed plausible by the decision maker.

A robust MDP starts in state $x \in \mathcal{X}$ with known probability $q(x)$. The decision maker can then select any action $a \in \mathcal{A}$, upon which an immediate reward of $c(x, a)$ is earned and the MDP transitions to state $x' \in \mathcal{X}$ with probability $p(x'|x, a)$, where p can be any element of the ambiguity set \mathcal{P} . The process then continues in the same fashion, governed by the same transition kernel p , for an infinite length of time. The decision maker wishes to determine a policy $\pi : \mathcal{X} \rightarrow \mathcal{A}$, which declares for each state $x \in \mathcal{X}$ which action $a \in \mathcal{A}$ is to be taken in state x , that maximizes the worst-case expected total discounted reward:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot c(x_t, a_t) \mid x_0 \sim q \right]$$

Here, Π denotes the set of all deterministic, memoryless policies $\pi : \mathcal{X} \rightarrow \mathcal{A}$, $\{(x_t, a_t)\}_{t=0}^{\infty}$ is the stochastic process induced by the initial probabilities q , the transition probabilities p and the policy π , and \mathbb{E} is the expectation operator with respect to this process.

We next construct a robust MDP for the unconstrained robust assortment optimization problem under the MC model.

DEFINITION 3 (ROBUST MDP REPRESENTATION). For a given MC model, we define the robust MDP $(\mathcal{X}, \mathcal{A}, \mathcal{P}, q, c, \gamma)$ via the state space $\mathcal{X} = \mathcal{N}_0$, the action space $\mathcal{A} = \{\top, \perp\}$, the ambiguity set \mathcal{P} that contains all transition kernels p satisfying

$$p(x'|x, \top) = \begin{cases} 1 & \text{if } x' = x, \\ 0 & \text{otherwise,} \end{cases} \quad p(x'|x, \perp) = \begin{cases} \rho_{xx'}/\gamma & \text{if } x, x' \neq 0, \\ 1 - \sum_{x \in \mathcal{N}} \rho_{xx}/\gamma & \text{if } x' = 0, \\ 0 & \text{otherwise} \end{cases} \quad \forall x, x' \in \mathcal{X} \quad (6)$$

for some $\boldsymbol{\rho} \in \mathcal{U}$, the initial state probabilities $q(x) = \lambda_x$, $x \in \mathcal{X}$, the immediate rewards $c(x, \top) = (1 - \gamma)r_x$ and $c(x, \perp) = 0$, $x \in \mathcal{X}$, and $\gamma = 1 - \min \{\rho_{i_0} : \boldsymbol{\rho} \in \mathcal{U}, i \in \mathcal{N}\}$.

Intuitively, the states of the robust MDP describe the different purchase options $x \in \mathcal{N}_0$ of the customer, and the actions \top and \perp characterize the options of the decision maker to include (exclude) any of the products $x \in \mathcal{N}$ in/from the assortment. Note that the virtual product 0 is always available, and hence the selected action does not matter in state 0. If the customer enters a state whose associated purchase option is part of the assortment, then this state keeps generating revenues and is never left. Otherwise, the next purchase option considered by the customer is selected randomly according to some $\boldsymbol{\rho} \in \mathcal{U}$.

We next show that [Definition 3](#) indeed describes a valid robust MDP.

OBSERVATION 1. *The robust MDP from [Definition 3](#) is well-defined.*

We can now establish the equivalence between the unconstrained robust assortment optimization problem under the MC model and our robust MDP from [Definition 3](#).

THEOREM 3. *The unconstrained robust assortment optimization problem under the MC model is equivalent to the robust MDP from [Definition 3](#) in the following sense:*

- (i) *For every $S \in \mathcal{S}$, the worst-case expected total discounted reward of any $\pi_S \in \Pi$ satisfying $\pi_S(i) = \top$, $i \in S$, and $\pi_S(i) = \perp$, $i \in \mathcal{N} \setminus S$, coincides with the revenues $R^*(S)$.*
- (ii) *For every $\pi \in \Pi$, the revenues $R^*(S)$ of $S = \{i \in \mathcal{N} : \pi(i) = \top\}$ coincide with the worst-case expected total discounted reward of π .*

[Theorem 3](#) shows that there is a one-to-many relationship between the assortments $S \in \mathcal{S}$ of the assortment optimization problem and the policies $\pi_S : \mathcal{X} \rightarrow \mathcal{A}$ through the relation $i \in S \Leftrightarrow \pi_S(i) = \top$, $i \in \mathcal{N}$. The relationship is not one-to-one since the no-purchase option 0 is always available, irrespective of whether $\pi_S(0) = \top$ or $\pi_S(0) = \perp$. [Theorem 3](#) allows us to apply the rich arsenal of solution methods for robust MDPs to the unconstrained deterministic robust assortment optimization problem under the MC model with product-wise substitution sets (*cf.* [Section 5.3](#)). In the next subsection, we will leverage the established theory for robust MDPs to investigate under which conditions the decision maker may benefit from randomizing between multiple assortments, that is, when $R_{\text{rand}}^*(\mathcal{U}) > R_{\text{det}}^*(\mathcal{U})$.

REMARK 1 (NOMINAL ASSORTMENT OPTIMIZATION). If the uncertainty set of the robust assortment optimization problem is a singleton, say $\mathcal{U} = \{\boldsymbol{\rho}^0\}$, then the robust MDP from [Definition 3](#) reduces to a nominal MDP. In that case, the conclusions from [Observation 1](#) and [Theorem 3](#) continue to apply.

5.2. Randomization-Receptiveness of the Markov Chain Problem

We saw that under the MNL model, the unconstrained robust assortment optimization problem is randomization-proof. We now investigate whether this remains the case for the more general MC model. It turns out that while certain special cases of this more general problem are randomization-proof, the firm may generally benefit from randomization even in the unconstrained setting, and this carries over to the cardinality-constrained setting.

We first consider the special case of uncertainty sets with *product-wise substitution sets*,

$$\mathcal{U} = \left\{ \boldsymbol{\rho} \in \mathbb{R}_+^{(n+1) \times (n+1)} : \exists \boldsymbol{\rho}^0, \dots, \boldsymbol{\rho}^n \in \mathcal{U} \text{ such that } \rho_{ij} = \rho_{ij}^i \forall i, j \in \mathcal{N}_0 \right\}.$$

One readily verifies that this condition is equivalent to requiring that

$$\mathcal{U} = \bigtimes_{i \in \mathcal{N}_0} \mathcal{U}_i, \quad \text{where } \mathcal{U}_i = \left\{ \boldsymbol{\rho}_i = (\rho_{i0}, \dots, \rho_{in}) \in \mathbb{R}_+^{n+1} : \boldsymbol{\rho} \in \mathcal{U} \right\}.$$

Intuitively, the uncertainty set \mathcal{U} has product-wise substitution sets when knowledge of the uncertain substitution probabilities $\boldsymbol{\rho}_i = (\rho_{i0}, \dots, \rho_{in})$ for product $i \in \mathcal{N}_0$ does not allow the decision maker to infer anything about the uncertain substitution probabilities $\boldsymbol{\rho}_j = (\rho_{j0}, \dots, \rho_{jn})$ for any other product $j \in \mathcal{N}_0, j \neq i$, beyond the fact that $\boldsymbol{\rho}_j \in \mathcal{U}_j$. In this case, similar to the MNL model, randomization offers no benefits in the unconstrained problem.

THEOREM 4. *If the uncertainty set \mathcal{U} of the unconstrained robust assortment optimization problem under the MC model has product-wise substitution sets, the problem is randomization-proof.*

The proof of [Theorem 4](#) shows that for product-wise substitution sets, the robust MDP from [Definition 3](#) has an (x, a) -rectangular ambiguity set, which implies strong duality:

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right] = \min_{p \in \mathcal{P}} \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right].$$

[Corollary 1](#) then implies that the problem is randomization-proof.

[Désir et al. \(2023\)](#) derive a strong duality result for the unconstrained robust assortment optimization problem under the MC model with product-wise substitution sets (termed *row-wise uncertainty* in that work). In contrast to our result, which establishes a connection between the unconstrained assortment optimization problem and robust MDPs and subsequently leverages existing results for robust MDPs, [Désir et al. \(2023\)](#) prove strong duality *ab initio*.

While product-wise substitution sets may appear to constitute an intuitive choice for the uncertainty set \mathcal{U} , they are unlikely to arise as a result of a statistical estimation from historical data. Instead, a data-driven estimation approach based on a confidence region formed, for example, by a maximum likelihood estimation, would exhibit an asymptotically elliptical shape under which all rows ρ_i of the uncertain substitution matrix ρ are dependent (cf. Billingsley 1961). Under such general substitution sets, strong duality no longer holds for the associated robust MDPs (cf. Wieseemann et al. 2013), which opens up the possibility for the problem to be randomization-receptive. This does not automatically imply, however, that randomization is beneficial, as the next result shows.

PROPOSITION 1. *Any instance of the unconstrained robust MC problem with two products, irrespective of the geometry of the uncertainty set \mathcal{U} , is randomization-proof.*

Generally, however, the more complex choice behavior captured by the MC model allows the firm to benefit from randomization. Specifically, with three or more products, we can find instances where the unconstrained robust MC problem with general substitution sets is receptive to randomization.

OBSERVATION 2. *For any number $n \geq 3$ of products, the unconstrained MC problem is randomization-receptive.*

In the constrained MNL problem, we found the potential benefits of randomization to be unbounded. In the unconstrained MC problem, however, this is not the case, as the following results shows.

OBSERVATION 3. *Under the unconstrained MC problem with general substitution sets, the benefits of randomization are bounded from above by*

$$\frac{R_{\text{rand}}^*(\mathcal{U})}{R_{\text{det}}^*(\mathcal{U})} \leq \frac{\max\{r_i : i \in \mathcal{N}\}}{\min\{r_i : i \in \mathcal{N}\}}.$$

The intuition underlying [Observation 3](#) is as follows. Imagine first that the customers do not exercise the outside option. In that case, the benefits of randomization cannot exceed the bound of [Observation 3](#) since in the unconstrained problem, the assortment containing all products \mathcal{N} is always feasible and generates revenues of at least $\min\{r_i : i \in \mathcal{N}\}$, while no randomized assortment strategy can generate more revenues than $\max\{r_i : i \in \mathcal{N}\}$. Since the probability of a customer exercising the outside strategy under assortment T is at most

as large as under any assortment $S \subseteq T$, one can show that the bound of [Observation 3](#) carries over to the case where customers can exercise the outside option.

We now turn to the cardinality-constrained problem. The next result confirms that similar to the MNL problem, the cardinality-constrained MC problem is randomization-receptive under product-wise and general substitution sets. Moreover, the benefits of randomization are again unbounded, and any optimal randomization strategy can be complex.

THEOREM 5. *For any number of products $n \geq 3$ and any restriction $|S| \leq C$, $C \in \{1, \dots, n-2\}$, there are instances of the cardinality-constrained robust MC problem with product-wise substitution sets where*

- $R_{\text{det}}^*(\mathcal{U}) = 0$ while $R_{\text{rand}}^*(\mathcal{U}) > 0$.
- the unique optimal randomized assortment strategy places equal (positive) probability on each assortment S satisfying $|S| = C$ and zero probability on all other assortments.

5.3. Solving the Randomized Markov Chain Problem

To solve the unconstrained and constrained randomized assortment optimization problem under the MC model with general substitution sets, we stipulate that the uncertainty set $\mathcal{U} \subseteq \mathbb{R}_+^{(n+1) \times (n+1)}$ is a finite set of substitution matrices. In the special case of product-wise substitution sets, we furthermore stipulate that $\mathcal{U} = \times_{i \in \mathcal{N}_0} \mathcal{U}_i$ for finite sets $\mathcal{U}_i \subseteq \mathbb{R}^{n+1}$ of substitution vectors for each product $i \in \mathcal{N}_0$. Popular examples of such uncertainty sets include projections of factor model uncertainty sets onto the cross-product of $n+1$ probability simplices $\Delta(\mathcal{N}_0)$. We present in [Appendix A](#) a two-layer column generation framework to solve the corresponding assortment optimization problems.

6. Preference Ranking Model

The PR model (see, *e.g.*, [Bertsimas and Mišić 2019](#)) is parameterized by K bijective preference rankings $\sigma_k : \mathcal{N}_0 \rightarrow \{1, \dots, n+1\}$, $k \in \mathcal{K} = \{1, \dots, K\}$, with occurrence probabilities $\lambda_k \in \mathbb{R}_+$ satisfying $\mathbf{e}^\top \boldsymbol{\lambda} = 1$. The probability that a customer is characterized by the preference ranking σ_k , $k \in \mathcal{K}$, is λ_k . Faced with the assortment S as well as the no-purchase option 0, such a customer purchases the product $i \in S \cup \{0\}$ that has the smallest rank in σ_k , that is, $i \in \arg \min \{\sigma_k(j) : j \in S \cup \{0\}\}$. Hence, the expected revenues of the assortment $S \in \mathcal{S}$ under the PR model amount to

$$R(S, \boldsymbol{\lambda}) = \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S), \tag{7}$$

where $R_k(S) = r_i$ for the unique $i \in S \cup \{0\}$ that satisfies $\sigma_k(i) < \sigma_k(j)$ for all $j \in S \cup \{0\}$, $j \neq i$. In the corresponding robust assortment optimization problem, we assume that the vector of occurrence probabilities $\boldsymbol{\lambda}$ is only known to be contained in a compact ambiguity set $\mathcal{U} = \Lambda \subseteq \{\boldsymbol{\lambda} \in \mathbb{R}_+^K : \mathbf{e}^\top \boldsymbol{\lambda} = 1\}$, and we seek to maximize the worst-case expected revenues.

Section 6.1 studies the potential benefits of randomization under the robust PR problem, and Section 6.2 discusses the solution of the (un-)constrained robust assortment optimization problem under the PR model.

6.1. Randomization-Receptiveness of the Preference Ranking Problem

Similar to the MC problem under general uncertainty sets, the robust assortment optimization problem under the PR model is randomization-receptive even in the absence of any further constraints (such as cardinality constraints). In fact, randomization can be beneficial for any number of products, as we show next.

OBSERVATION 4. *For any number of products $n \geq 2$, the unconstrained as well as the constrained PR problem are randomization-receptive.*

Contrary to the constrained MNL and MC problems (cf. Theorems 2 and 5), however, and similar to the unconstrained MC problem with general substitution sets (cf. Observation 3), the benefits of randomization are then bounded.

OBSERVATION 5. *For any instance of the unconstrained robust PR problem with $n = 2$ products, the benefits of randomization are bounded from above by $R_{\text{rand}}^*(\mathcal{U}) / R_{\text{det}}^*(\mathcal{U}) \leq 2$. Moreover, for any instance the benefits of randomization are bounded from above by*

$$\frac{R_{\text{rand}}^*(\mathcal{U})}{R_{\text{det}}^*(\mathcal{U})} \leq \frac{\max\{r_i : i \in \mathcal{N}\}}{\min\{r_i : i \in \mathcal{N}\}}.$$

When we impose a cardinality constraint $|S| \leq C$ on the size of the offered assortment, similar to the cardinality-constrained MNL and MC problems (cf. Theorems 2 and 5), the potential benefits of randomization in the cardinality-constrained PR problem can be arbitrarily large.

THEOREM 6. *For any number of products $n \geq 2$ and any restriction $|S| \leq C$, $C \in \{1, \dots, n-1\}$, there are instances of the cardinality-constrained robust PR problem where $R_{\text{det}}^*(\mathcal{U}) = 0$ while $R_{\text{rand}}^*(\mathcal{U}) > 0$.*

We have seen that in the cardinality-constrained MNL and MC problems, optimal policies may require randomization between an exponentially large number of assortments (cf. [Theorem 2](#) and [Theorem 5](#)). By contrast, under the PR model (whether constrained or not), there is always a parsimonious optimal randomization strategy.

THEOREM 7. *Under the PR model (with arbitrary constraints on the set of admissible assortments), there always exists an optimal randomization strategy which places strictly positive weight on no more than $K + 1$ assortments. Moreover, when the uncertainty set is an affine map of a polyhedral subset of \mathbb{R}^m , there exists an optimal randomization strategy which places strictly positive weight on no more than $m + 1$ assortments.*

6.2. Solving the Randomized Preference Ranking Problem

[Theorem 7](#) allows us to formulate the unconstrained and constrained randomized assortment optimization problem under the PR model as robust K -adaptability problems that determine $K + 1$ assortments and their randomization weights ([Bertsimas and Caramanis 2010](#), [Hanasusanto et al. 2015](#)). While the resulting problems can be expressed as mixed-integer linear programs of compact size that are amenable to solution via standard solvers, our numerical experiments indicate that this approach does not scale to problems of interesting size. Instead, we develop an exact and a heuristic column generation scheme to compute randomization strategies for the PR model. To this end, we assume that the ambiguity set Λ is a polyhedral set of the form

$$\Lambda = \{ \boldsymbol{\lambda} = \mathbf{F}\boldsymbol{\xi} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}, \boldsymbol{\xi} \in \mathbb{R}_+^m \}, \quad (8)$$

where $\mathbf{F} \in \mathbb{R}^{K \times m}$, $\mathbf{A} \in \mathbb{R}^{l \times m}$ and $\mathbf{b} \in \mathbb{R}^l$. We list below two popular choices of such sets.

(i) **Norm ambiguity sets.** For a nominal ranking prevalence vector $\boldsymbol{\lambda}^0 \in \mathbb{R}_+^K$ and a radius $\theta \in \mathbb{R}_+$, we define the ambiguity set

$$\Lambda = \left\{ \boldsymbol{\lambda} : \mathbf{e}^\top \boldsymbol{\lambda} = 1, \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^0\|_p \leq \theta, \boldsymbol{\lambda} \in \mathbb{R}_+^K \right\},$$

where $p \in \{1, \infty\}$. Norm ambiguity sets hedge against all perturbations of the ranking prevalence vector $\boldsymbol{\lambda}$ that are contained in a p -ball of radius θ around the nominal vector $\boldsymbol{\lambda}^0$. In particular, the choice $p = 1$ recovers the total variation distance, a popular ϕ -divergence ([Bayraksan and Love 2015](#)), which allows us to choose the radius θ based on statistical bounds.

(ii) **Approximate ellipsoidal ambiguity sets.** For a nominal ranking prevalence vector $\boldsymbol{\lambda}^0 \in \mathbb{R}^K$ and a symmetric and positive definite matrix $\mathbf{P} \in \mathbb{R}^{K \times K}$, we set

$$\Lambda = \left\{ \boldsymbol{\lambda} : \mathbf{e}^\top \boldsymbol{\lambda} = 1, \boldsymbol{\lambda} = \boldsymbol{\lambda}^0 + \mathbf{P}\boldsymbol{\xi}, \|\boldsymbol{\xi}\|_1 \leq \sqrt{K}, \|\boldsymbol{\xi}\|_\infty \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^K, \boldsymbol{\xi} \in \mathbb{R}^K \right\}.$$

Note that $\{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_2 \leq 1\} \subseteq \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_1 \leq \sqrt{K}\} \cap \{\boldsymbol{\xi} \in \mathbb{R}^K : \|\boldsymbol{\xi}\|_\infty \leq 1\}$, and thus the ambiguity set constitutes an outer (conservative) approximation of an ellipsoid with center $\boldsymbol{\lambda}^0$ and semi-axes defined by \mathbf{P} . Ellipsoidal ambiguity sets recover the Pearson χ^2 -divergence, another popular ϕ -divergence (Bayraksan and Love 2015), and they emerge asymptotically as confidence regions of a maximum likelihood estimation.

The randomized robust assortment optimization problem under the PR model can be formulated as the following robust linear program:

$$\begin{aligned} & \text{maximize} && \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} \sum_{k \in \mathcal{K}} p_S \cdot \lambda_k \cdot R_k(S) \\ & \text{subject to} && \sum_{S \in \mathcal{S}} p_S = 1 \\ & && p_S \geq 0, S \in \mathcal{S} \end{aligned} \tag{9}$$

Problem (9) is computationally challenging as it involves exponentially many decision variables $p_S, S \in \mathcal{S}$. For our solution schemes, it is useful to consider the dual of problem (9):

$$\begin{aligned} & \text{minimize} && \max_{S \in \mathcal{S}} \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S) \\ & \text{subject to} && \boldsymbol{\lambda} \in \Lambda \end{aligned} \tag{10}$$

Strong duality between (9) and (10) holds since (9) is feasible by construction. Although problem (10) contains only polynomially many decision variables, the optimization problem embedded in its objective function maximizes over a combinatorial set $S \in \mathcal{S}$. Thus, neither the primal problem (9) nor the dual problem (10) is amenable to a solution with an off-the-shelf solver. We present in Appendix A.3 a column generation scheme for problem (9) that can be interpreted as a cutting plane approach for problem (10).

7. Numerical Results

This section presents numerical results for both synthetic and real-world data. The results in Sections 7.1–7.3 aim to elucidate when randomization can help to improve the worst-case expected revenues on synthetic data sets under the cardinality-constrained MNL problem, the MC problem, and the unconstrained PR problem. Moreover, for the MNL problem,

we demonstrate how the improvement of the worst-case expected revenues can translate into improvements of the out-of-sample expected revenues in a data-driven setting. We also investigate the computational price to be paid for exact and heuristic solutions to the deterministic and randomized robust assortment optimization problems. Finally, in Section 7.4, we demonstrate the value of randomization on a real-world data set. All solution schemes are implemented in C++ and run on Intel Xeon 2.20GHz cluster nodes with 16 GB dedicated main memory in four-core mode. Our data sets and detailed results, together with the source codes of all our algorithms, can be found online.²

7.1. The Cardinality-Constrained Multinomial Logit Problem

We first consider the cardinality-constrained MNL model where the product prices r_i are selected uniformly at random from the interval $[0, 10]$. We use a budget uncertainty set (cf. Section 4.2) where the lower and upper product valuations \underline{v}_i and \bar{v}_i , $i \in \mathcal{N}$, are chosen uniformly at random from the intervals $[0, 4]$ and $[6, 10]$, respectively, whereas the valuation of the no-purchase option is fixed at $v_0 = 5$.

Table 3 presents results for $n = 20$ products where the uncertainty budgets Γ (rows) and the assortment cardinalities $|S| \leq C$ (columns) are set to various percentages of n . For each table entry, the first (upper) value denotes the percentage of 250 randomly generated instances in which the optimal randomized assortment outperformed the optimal deterministic robust assortment in terms of worst-case expected revenues, while the second (lower) value reports the average outperformance on those instances. The table shows that the benefits of randomization are most significant when Γ is close to C and both quantities are small relative to the number of products. The latter is intuitive as $C = n$ recovers the randomization-proof unconstrained robust MNL problem while $\Gamma = n$ recovers the randomization-proof nominal MNL problem under the valuations $\mathbf{v}^0 = \underline{\mathbf{v}}$ (cf. Example 3.2 of Rusmevichientong and Topaloglu 2012).

We next investigate the benefits of randomization when the number n of products varies. To this end, we select $C = \Gamma = \lfloor \frac{1}{2}\sqrt{n} \rfloor$, which is in line with the central limit theorem-type uncertainty budget sets proposed by Bandi and Bertsimas (2012). The ‘randomized exact objective’ column of Table 4 shows that the benefits of randomization (measured here in terms of the average outperformance over the deterministic robust assortment over

² www.doc.ic.ac.uk/~wwiesema/assortment_opt.zip

		C							
		5%	10%	15%	20%	25%	30%	50%	75%
Γ	5%	100.00%	98.00%	83.60%	52.40%	20.40%	6.80%	0.40%	0.40%
		33.50%	4.17%	0.97%	0.35%	0.20%	0.10%	0.00%	0.00%
	10%	96.40%	100.00%	96.40%	79.60%	48.40%	17.20%	0.00%	0.00%
		20.20%	19.99%	5.09%	1.60%	0.54%	0.24%	0.00%	0.00%
	15%	74.00%	95.20%	99.60%	92.00%	68.00%	35.60%	0.00%	0.00%
		13.94%	12.18%	13.12%	4.48%	1.50%	0.54%	0.00%	0.00%
	20%	50.00%	80.00%	92.40%	97.20%	82.40%	56.80%	0.00%	0.00%
		9.49%	7.19%	7.63%	8.67%	3.42%	1.12%	0.00%	0.00%
	25%	26.40%	50.40%	70.40%	84.40%	90.00%	70.40%	1.60%	0.80%
		6.87%	4.84%	4.56%	4.92%	5.72%	2.36%	0.04%	0.01%
30%	11.20%	27.60%	41.60%	61.20%	74.80%	80.00%	1.20%	0.40%	
	5.27%	3.29%	2.85%	2.68%	3.00%	3.46%	0.15%	0.01%	
50%	0.00%	0.00%	0.00%	0.40%	1.20%	1.60%	8.40%	0.40%	
	0.00%	0.00%	0.00%	0.26%	0.28%	0.26%	1.07%	0.01%	
75%	0.00%	0.00%	0.00%	0.00%	0.40%	0.40%	0.80%	0.40%	
	0.00%	0.00%	0.00%	0.00%	0.01%	0.01%	0.00%	0.01%	

Table 3 Benefits of randomization in the cardinality-constrained MNL problem. In the table, the rows (columns) correspond to different uncertainty budgets Γ (assortment cardinalities C).

250 random problem instances), while decreasing with problem size, is significant for all considered instance sizes. Interestingly, our exact column generation scheme for the randomized problem is actually faster than the cutting plane technique that we implemented for the deterministic robust problem, as the columns ‘deterministic exact runtime’ and ‘randomized exact runtime’ in Table 4 reveal. This is caused by the fact that the cutting plane technique requires significantly more iterations (median 14 for $n = 25$ and median 47 for $n = 50$, for example) than the column generation scheme (median 5 for $n = 25$ and median 7 for $n = 50$, with 2-5 primal and dual iterations per main iteration). Since each (main) iteration adds a fractional linear term to the problems that requires an individual reformulation resulting in additional auxiliary variables and big-M constraints, the number of (main) iterations is a key performance indicator for both algorithms.

The columns ‘deterministic heuristic’ and ‘randomized heuristic’, finally, report the runtimes and objective values of the ADXOpt heuristic for the deterministic and randomized robust problems, respectively. In both cases, the objective values are again measured relative to the worst-case expected revenues of the exact deterministic robust problem. We see that for both problems, the ADXOpt heuristic performs very well, particularly on larger instance sizes. For the deterministic robust problem, the optimality gaps approach 0% as the instance sizes grow, while for the randomized robust problem, the outperformance of the ADXOpt heuristic approaches the outperformance of the exact solution approach. We

thus conclude that one should solve the deterministic and randomized robust problems exactly for small problem sizes, in which case the cutting plane and column generation schemes are fast, while one may want to resort to heuristic solutions for larger problem sizes, where the optimality gap of ADXOpt decreases rapidly.

n	deterministic			randomized			
	exact runtime	heuristic objective	runtime	exact objective	runtime	heuristic objective	runtime
5	0.02	-6.80%	0.00s	37.95%	0.01s	24.97%	0.00s
10	0.05	-1.83%	0.00s	35.15%	0.03s	32.03%	0.00s
15	0.13	-0.96%	0.00s	35.39%	0.08s	34.01%	0.00s
20	2.30	-0.84%	0.00s	21.30%	1.80s	20.11%	0.03s
25	6.94	-0.41%	0.01s	22.72%	3.58s	21.96%	0.05s
30	18.38	-0.38%	0.01s	21.99%	5.88s	21.15%	0.09s
35	27.18	-0.40%	0.02s	21.63%	8.10s	21.04%	0.12s
40	299.18	-0.23%	0.05s	9.41%	64.21s	8.58%	0.76s
45	592.28	-0.12%	0.07s	7.29%	97.56s	6.71%	1.09s
50	937.29	-0.11%	0.09s	6.09%	171.07s	5.66%	1.57s

Table 4 Exact and heuristic solutions for the cardinality-constrained MNL problem. All percentages are reported relative to the optimal solutions of the deterministic robust problem.

We close this section with a data-driven experiment. To this end, we fix $n = 10$ products and consider different cardinalities $C \in \{1, 2, 3, 4\}$ for the admissible assortments. The product prices r_i are generated randomly by the same procedure as before. The true customer valuations are unknown and satisfy $v_i^0 = e^{\beta_i^0}$, where β_i^0 is drawn uniformly at random from -3 to 3 , $i \in \mathcal{N}$; we fix $v_0^0 = e^0$ and assume that this quantity is known. We assume that 5, 10, \dots , 95 historical samples of random assortments of cardinality C are available, together with the purchase choice that each customer made. In the nominal model, we then estimate the valuations $\hat{\mathbf{v}}$ from the historical data using a maximum likelihood estimation, and we solve the resulting nominal assortment optimization problem. In the deterministic and randomized robust models, we employ the budget uncertainty set where both $\Gamma \in \{0, \dots, n\}$ and $(\underline{\mathbf{v}}, \bar{\mathbf{v}}) = ([1 - \gamma]\hat{\mathbf{v}}, [1 + \gamma]\hat{\mathbf{v}})$, $\gamma \in \{0, 0.025, \dots, 0.5\}$ are selected using 7-fold cross-validation on the available historical data. Since the validation data reports the customer choices for assortments that differ from the assortment computed using the training data, each iteration of our cross-validation proceeds in two steps. We first use the validation set to perform a maximum likelihood estimation of the valuations. We subsequently use the nominal MNL model associated with these estimated valuations to approximate the out-of-sample revenues of the assortment computed using the training




Figure 2 Data-driven experiment for the MNL problem with cardinality constraint $C = 1, 2, 3, 4$, from left to right. All optimality gaps are reported relative to the expected revenues of the clairvoyant model that knows the true customer valuations.

data. We then compare the actual out-of-sample expected revenues of the three models under newly generated data using the true valuations \mathbf{v}^0 .

For each cardinality $C \in \{1, 2, 3, 4\}$ and sample size 5, 10, \dots , 95, Figure 2 reports the average optimality gaps (over 100 randomly generated instances) of the out-of-sample expected revenues of each model relative to the expected revenues of the nominal model under the true valuations. The figure shows that these optimality gaps decrease with sample size (*i.e.*, from left to right in each graph) as well as the cardinality of the assortment (*i.e.*, from the leftmost to the rightmost graph). This is expected as in both cases, the estimation problem can rely on more data and hence becomes easier. In all cases, the randomized robust model outperforms the deterministic robust model, which in turn outperforms the nominal model. We emphasize that this is not *a priori* obvious as the deterministic and randomized robust models use the available historical data for both estimation and parameter selection, and hence have less data than the nominal model to estimate $\hat{\mathbf{v}}$. Interestingly, for small cardinalities—where the estimation problem is most challenging—the randomized robust model significantly outperforms both the nominal and the deterministic robust model. For larger cardinalities, the performance of the deterministic and randomized robust models are becoming more similar. It is noteworthy, however, that the randomized robust model is never performing worse than the deterministic robust model (in terms of average optimality gap), while it is at the same time easier to solve (*cf.* Table 4).

7.2. The Markov Chain Problem

We next consider the MC model where the product prices are again selected uniformly at random. We study the randomization-receptive cases of this problem identified in Section 5, namely the cardinality-constrained problem under product-wise substitution sets, and both the unconstrained and the constrained problem under general substitution sets. We employ the factor model uncertainty sets outlined in Section 5.3 with K factors leading to uncertainty sets with 2^K realizations under general substitution sets and $(2^K)^n$ realizations under product-wise substitution sets, respectively.

We first study the MC model under general substitution sets. To this end, we compare the deterministic and randomized robust solutions for $K \in \{5, 8, 10, 12\}$ factors, $n \in \{20, 30, \dots, 90\}$ products and a cardinality constraint, when used, of $|S| \leq C = 5$. We determine the optimal randomized assortments using the method described in Appendix A.2. The left-hand side of Table 5 shows the results for the unconstrained problem. The rows and columns in the table correspond to the numbers of products n and the number of factors K of the uncertainty set, respectively. For each entry, the two numbers again denote the fraction of 250 randomly generated instances that were randomization-receptive (top) and the average outperformance of the randomized robust solution over the deterministic robust one on these instances (bottom). The table shows that the benefits of randomization are highest when the number of products is relatively low while the number of factors is high. The right side of the same table shows the corresponding results for the cardinality-constrained problem. Introducing the cardinality constraint significantly increases both the likelihood and the magnitude of randomization gains. The cardinality-constrained problem is, however, also more difficult to solve, and we were unable to obtain solutions for $n \geq 90$ when $K = 12$. Tables 12 and 13 in Appendix F compare the solution times for the deterministic robust problem with those of the randomized robust problem.

We next consider the cardinality-constrained MC problem with product-wise substitution sets. This problem variant turns out to be substantially harder to solve due to the large cardinality $(2^K)^n$ of the uncertainty set \mathcal{U} . We therefore fix $K = 2$ and limit the runtime of Step 2(b) in each iteration of our column generation scheme (see Appendix A.2) to 100 seconds. Table 6 reports the relative outperformance of our (heuristic, due to the premature termination) randomized assortments over the best deterministic assortments. The outperformance is reported in terms of lower (top entry) and upper (bottom entry)

		K						K			
		5	8	10	12			5	8	10	12
n	20	93.60%	97.60%	98.80%	99.60%	20	99.20%	100.00%	100.00%	100.00%	
		0.91%	1.29%	1.35%	1.48%		4.51%	7.53%	8.70%	9.20%	
	30	95.60%	98.80%	99.20%	98.80%	30	100.00%	100.00%	100.00%	100.00%	
		0.65%	0.81%	0.96%	1.05%		6.35%	9.93%	12.01%	13.48%	
	40	95.20%	96.76%	98.80%	99.60%	40	100.00%	100.00%	100.00%	100.00%	
		0.45%	0.60%	0.63%	0.73%		7.85%	12.52%	15.20%	16.94%	
	50	93.20%	99.60%	97.20%	99.20%	50	100.00%	100.00%	100.00%	100.00%	
		0.37%	0.44%	0.47%	0.55%		9.17%	14.82%	17.45%	17.80%	
	60	95.20%	97.60%	98.00%	98.80%	60	100.00%	100.00%	100.00%	100.00%	
		0.29%	0.32%	0.37%	0.41%		9.99%	15.90%	18.22%	18.27%	
	70	95.20%	98.80%	98.80%	99.19%	70	100.00%	100.00%	100.00%	100.00%	
		0.22%	0.26%	0.32%	0.32%		11.40%	16.60%	17.82%	16.93%	
80	94.40%	98.00%	100.00%	99.59%	80	100.00%	100.00%	100.00%	100.00%		
	0.20%	0.23%	0.27%	0.27%		11.92%	17.56%	16.79%	17.70%		
90	97.60%	99.20%	100.00%	99.02%	90	100.00%	100.00%	100.00%	n/a		
	0.17%	0.17%	0.22%	0.22%		12.94%	17.15%	16.19%	n/a		
100	97.20%	98.80%	98.33%	100.00%	100	100.00%	100.00%	100.00%	n/a		
	0.17%	0.17%	0.18%	0.21%		13.46%	16.63%	15.24%	n/a		

Table 5 Benefits of randomization under the MC problem with general substitution sets. The rows and columns correspond to the number of products n and the number of factors K , respectively.

Left table: unconstrained problem. Right table: cardinality constraint $|S| \leq C = 5$.

		n					
		10	12	14	16	18	20
lower bound (%)		1.15	2.24	4.70	6.94	7.11	2.45
upper bound (%)		1.15	2.24	4.70	6.99	8.61	12.24

Table 6 Outperformance of the randomized assortments over the deterministic ones under the MC model with product-wise substitution sets and the cardinality constraint $|S| \leq C = 5$.

bounds since the worst-case evaluation problem is terminated prematurely. The results represent averages over 100 randomly generated instances. While it is difficult to draw definite conclusions due to the widening gap of the bounds, the benefits of randomized assortments appear to grow with the number of products.

7.3. The Preference Ranking Problem

We next consider the PR model where the product prices are again selected uniformly at random from the interval $[0, 10]$. We use a 1-norm ambiguity set (*cf.* Section 6.2) where the nominal ranking prevalence vector λ^0 is selected uniformly at random from the K -dimensional probability simplex.

In our first experiment, we fix the number of products to $n = 20$ and vary the radius $\theta \in \{0.25, 0.5, \dots, 2\}$ of the ambiguity set as well as the number of preference rankings $K \in \{10, 100, 200, 500, 1,000\}$. Table 7 presents the percentages of instances (across 250 randomly generated instances) where the randomized robust model improved upon the

deterministic robust one in terms of worst-case expected revenues (first number) as well as the average relative outperformance across those instances (second number). The table shows that the benefits of randomization are most pronounced when the ambiguity is large but not covering the entire probability simplex (which is the case when $\theta = 2$). This is intuitive since under full ambiguity, nature can choose a Dirac distribution that places all probability mass on the least favorable (assortment-dependent) preference ranking, in which case randomization does not help. Likewise, if the ambiguity is small, we approach a non-robust problem where randomization does not help either.

		K				
		10	100	200	500	1,000
θ	0.00	0.00%	0.00%	0.00%	0.00%	0.00%
		0.00%	0.00%	0.00%	0.00%	0.00%
	0.25	3.00%	13.30%	13.30%	15.45%	12.88%
		0.60%	0.12%	0.06%	0.06%	0.03%
	0.50	30.04%	60.09%	58.80%	50.64%	62.23%
		1.21%	0.29%	0.21%	0.19%	0.23%
	0.75	25.75%	78.11%	80.69%	83.26%	78.97%
		0.63%	0.55%	0.51%	0.49%	0.46%
	1.00	33.48%	88.41%	94.85%	93.56%	94.85%
		1.99%	1.06%	1.00%	1.08%	0.96%
	1.25	29.18%	95.71%	97.00%	96.14%	97.85%
		3.08%	2.37%	2.50%	1.91%	1.84%
	1.50	19.74%	98.71%	98.71%	100.00%	100.00%
		2.63%	3.91%	3.47%	3.38%	3.04%
	1.75	0.00%	93.56%	99.57%	100.00%	99.57%
		0.00%	10.81%	9.22%	6.49%	7.17%
	2.00	0.00%	0.00%	0.00%	0.00%	0.00%
		0.00%	0.00%	0.00%	0.00%	0.00%

Table 7 Benefits of randomization in the PR problem. The rows (columns) correspond to different ambiguity radii θ (numbers of preference rankings K).

We now investigate the benefits of randomization when the number n of products varies. To this end, Table 8 fixes the radius of ambiguity to $\theta = 1.5$ and reports the outperformance of the exact (column ‘randomized exact objective’) and heuristic (column ‘randomized heuristic objective’) solutions to our randomized robust model relative to the optimal solution to the deterministic robust model. The table shows that while the benefits of randomization decrease with the number of products n , they remain non-trivial for all considered problem sizes. Moreover, the performance of the ADXOpt heuristic is more or less independent of the number of products n as well as improving with the number of preference rankings K .

Table 8 also reports the runtimes of the mixed-integer programs for the deterministic robust model (column ‘deterministic exact runtime’) as well as our column generation scheme (column ‘randomized exact runtime’) and the ADXOpt heuristic (column ‘randomized heuristic runtime’) for the randomized robust model. The table shows that the randomized exact problem is significantly more challenging to solve than deterministic one. Indeed, while the exact deterministic robust problem can be solved as a monolithic MILP, the exact randomized robust problem is solved in a cutting plane fashion and thus suffers from similar scalability issues as the cutting plane scheme for the deterministic robust MNL problem (*cf.* Table 4). The number of cutting plane iterations grows with the problem size (median 6 for $n = 20$, $K = 10$ vs. median 49 for $n = 50$, $K = 1,000$), and each subproblem of the algorithm requires the solution of a MILP whose number of decision variables scales in K . The ADXOpt heuristic, on the other hand, provides close-to-optimal solutions while requiring substantially less computational resources.

We have also attempted to solve the randomized robust assortment optimization problem under the PR model as a K -adaptability problem (*cf.* Section 6.2). It turns out, however, that the resulting mixed-integer optimization problems grow quickly in size and become difficult to solve for instances with more than 10 products. We therefore do not report computational results for this solution scheme.

7.4. Experiments with Real-World Data

In this section, we demonstrate the value of randomized strategies on real data. To this end, we use the data set of Bertsimas and Mišić (2017, 2019), which is derived from a field study on consumer preferences for hypothetical Timbuk2 laptop bag designs by Toubia et al. (2003) (see also Belloni et al. 2008). The data set contains 330 preference rankings for 3,584 bag designs, as well as the revenues associated with each design.³ Bertsimas and Mišić (2017) use the data to study the value of accounting for uncertainty in deterministic robust assortment optimization under the PR model (among others). Here, we extend their approach to the randomized robust assortment optimization problem.

Bertsimas and Mišić (2017) solve the nominal and deterministic robust problems heuristically for the entire data set. By contrast, we use the approach described in Section 6.2 to solve the nominal as well as the deterministic and randomized robust problems exactly

³ The data set is available at <https://github.com/vvmisic/optimalPLD>.

n	K	deterministic	randomized			
		exact runtime	exact objective	exact runtime	heuristic objective	heuristic runtime
10	10	0.01s	0.54%	0.00s	0.51%	0.00s
10	100	0.05s	4.03%	0.12s	3.64%	0.03s
10	200	0.10s	4.95%	0.32s	4.10%	0.07s
10	500	0.41s	3.11%	0.90s	2.74%	0.18s
10	1000	0.68s	3.44%	2.57s	2.92%	0.49s
20	10	0.01s	0.44%	0.01s	-0.15%	0.00s
20	100	0.16s	3.83%	0.81s	3.16%	0.08s
20	200	0.34s	5.32%	2.48s	4.32%	0.18s
20	500	1.41s	3.38%	7.49s	2.94%	0.54s
20	1000	5.02s	2.99%	20.59s	2.61%	1.09s
30	10	0.01s	0.41%	0.02s	0.11%	0.00s
30	100	0.33s	3.82%	3.11s	2.89%	0.16s
30	200	0.78s	3.69%	7.38s	3.13%	0.40s
30	500	4.43s	2.97%	31.24s	2.61%	1.03s
30	1000	16.30s	2.91%	102.33s	2.62%	2.16s
40	10	0.03s	0.14%	0.03s	-0.01%	0.00s
40	100	0.49s	4.83%	6.43s	3.81%	0.33s
40	200	1.66s	3.23%	22.06s	2.52%	0.71s
40	500	8.73s	2.40%	82.31s	2.13%	1.81s
40	1000	43.61s	2.32%	362.83s	2.07%	3.83s
50	10	0.03s	0.08%	0.03s	-0.08%	0.01s
50	100	0.92s	2.17%	8.19s	1.74%	0.45s
50	200	2.56s	3.95%	53.55s	3.20%	1.22s
50	500	17.28s	2.71%	219.30s	2.33%	2.91s
50	1000	75.45s	2.46%	1011.82s	2.24%	6.99s

Table 8 Exact and heuristic solutions for the PR problem. All percentages are reported relative to the optimal solutions of the deterministic robust problem.

for a subset of available products. Specifically, we pick the $n = 300$ products generating the highest average consumer utilities in the data set and trim the preference ranking lists to only include these products and the no-purchase option. Following [Bertsimas and Mišić \(2017\)](#), we assume that each customer type is equally probable, that is, $\lambda_k = 1/K$ for all $k \in \mathcal{K}$. For the deterministic and randomized robust problems, we employ a 1-norm ambiguity set with varying radius θ (*cf.* Section 6.2).

Following [Bertsimas and Mišić \(2017\)](#), we use two metrics to quantify the value of the deterministic and randomized robust approaches: the worst-case loss and the relative improvement. The worst-case loss (WCL) measures the largest possible loss of the optimal nominal assortment S^N over the uncertainty set \mathcal{U} :

$$\text{WCL} = 100\% \cdot \frac{R_{\text{nom}}^* - R^*(S^N)}{R_{\text{nom}}^*},$$

where R_{nom}^* is the expected revenues of S^N under the nominal occurrence probabilities λ_0 and $R^*(S^N)$ denotes the worst-case expected revenues of S^N over the uncertainty set \mathcal{U} .

	WCL(%)	RI ^D (%)	RI ^R (%)	RI RD (%)
θ 0.1	6.07	0.15	0.15	0.00
0.2	12.14	0.34	0.35	0.01
0.3	18.22	0.56	0.78	0.21
0.4	24.09	0.76	2.74	1.96
0.5	29.88	0.88	6.62	5.69
0.6	35.68	1.02	11.73	10.60

Table 9 Benefits of randomization in the real-world case study over $n = 300$ products. The rows correspond to different radii θ of the uncertainty set.

The second metric is the relative improvement (RI) in worst-case expected revenues when using the deterministic or randomized robust approach instead of the nominal assortment:

$$\text{RI}^{\text{D}} = 100\% \cdot \frac{R_{\text{det}}^*(\mathcal{U}) - R^*(S^{\text{N}})}{R^*(S^{\text{N}})}$$

and

$$\text{RI}^{\text{R}} = 100\% \cdot \frac{R_{\text{rand}}^*(\mathcal{U}) - R^*(S^{\text{N}})}{R^*(S^{\text{N}})},$$

where $R_{\text{det}}^*(\mathcal{U})$ and $R_{\text{rand}}^*(\mathcal{U})$ denote the optimal values of the deterministic and randomized robust assortment optimization problems, respectively.

Table 9 reports the worst-case losses and relative improvements for different radii θ . The table shows that the worst-case loss from solving the nominal problem grows with the radius of the uncertainty set, as expected. The deterministic robust approach (column RI^D), consistently with Bertsimas and Mišić (2017), benefits the decision maker in terms of the relative improvement, and the magnitude of the improvement increases with the radius of the uncertainty set. The relative improvement achieved by the randomized robust strategy (RI^R) is significantly higher, however, demonstrating the additional benefit of randomization in hedging against uncertainty in the preference rankings. The difference in relative improvements between the two robust approaches is reported in the column RIRD, which shows that the additional gains of the randomized approach are substantial and grow with the size of the uncertainty set. The results thus indicate the potential promise of employing randomization in robust assortment optimization on real-world data.

8. Concluding Remarks

Our results call for more research into the importance of randomization for revenue management. We have seen that a firm can benefit from using randomization—instead of the standard approach of deterministically offering a single assortment—in the robust assortment optimization problem. However, it is not *a priori* obvious which problem settings may

benefit from randomization: two similar models, such as the MC model with product-wise and general uncertainty sets, may lead to starkly different conclusions. While our analysis suggests that more general versions of the assortment optimization problem tend to be more receptive to the benefits of randomization, more research is needed into this potential benefit under other choice models and constraints, as well as other revenue management problems. Also, all of the computational approaches developed in this paper study the static problem formulation where the decision maker does not update her knowledge about the uncertain problem parameters. As a next step, it would be instructive to study optimal dynamic randomization strategies where the decision maker learns about the uncertain problem parameters over time.

Acknowledgments

The first author gratefully acknowledges funding from the Imperial College President's PhD Scholarship programme. The last author gratefully acknowledges funding from the Engineering and Physical Sciences Research Council under the grants EP/R045518/1, EP/T024712/1 and EP/W003317/1. All three authors are indebted to the area editor (Gustavo Vulcano) as well as the anonymous associate editor and referees, whose insightful comments have led to substantial improvements of the manuscript.

References

- Aouad A, Farias V, Levi R. 2021. Assortment optimization under consider-then-choose choice models. *Management Science* **67**(6) 3368–3386.
- Aouad A, Farias V, Levi R, Segev D. 2018. The approximability of assortment optimization under ranking preferences. *Operations Research* **66**(6) 1661–1669.
- Bandi C, Bertsimas D. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming* **134** 23–70.
- Baron O, Milner J, Naseraldin H. 2011. Facility location: A robust optimization approach. *Production and Operations Management* **20**(5) 772–785.
- Bayraksan G, Love DK. 2015. Data-driven stochastic programming using phi-divergences. *INFORMS Tutorials in Operations Research* 1–19.
- Belloni A, Freund R, Selove M, Simester D. 2008. Optimizing product line designs: Efficient methods and comparisons. *Management Science* **54**(9) 1544–1552.
- Ben-Tal A, El Ghaoui L, Nemirovski A. 2009. *Robust Optimization*. Princeton University Press.
- Bertsimas D, Caramanis C. 2010. Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control* **55**(12) 2751–2766.
- Bertsimas D, Brown DB, Caramanis C. 2011. Theory and applications of robust optimization. *SIAM Review* **53**(3) 464–501.
- Bertsimas D, Mišić V. 2017. Robust product line design. *Operations Research* **65**(1) 19–37.
- Bertsimas D, Mišić V. 2019. Exact first-choice product line optimization. *Operations Research* **67**(3) 651–670.
- Bertsimas D, Nasrabadi E, Orlin JB. 2016. On the power of randomization in network interdiction. *Operations Research Letters* **44**(1) 114–120.
- Bertsimas D, Sim M. 2004. The price of robustness. *Operations Research* **52**(1) 35–53.
- Bertsimas D, Thiele A. 2006. A robust optimization approach to inventory theory. *Operations Research* **54**(1) 150–168.
- Billingsley P. 1961. *Statistical Inference for Markov Processes*. The University of Chicago Press.
- Birbil Şİ, Frenk J, Gromicho JA, Zhang S. 2009. The role of robust optimization in single-leg airline revenue management. *Management Science* **55**(1) 148–163.
- Bishop CM. 2006. *Pattern Recognition and Machine Learning*. Springer.
- Blanchet J, Gallego G, Goyal V. 2016. A Markov chain approximation to choice modeling. *Operations Research* **64**(4) 886–905.
- Bront JJM, Méndez-Díaz I, Vulcano G. 2009. A column generation algorithm for choice-based network revenue management. *Operations Research* **57**(3) 769–784.

- Bumpensanti P, Wang H. 2020. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science* **66**(7) 2993–3009.
- Şen A, Atamtürk A, Kaminsky P. 2018. A conic integer optimization approach to the constrained assortment problem under the mixed multinomial logit model. *Operations Research* **66**(4) 994–1003.
- Davis J, Gallego G, Topaloglu H. 2013. Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, Cornell University, New York.
- Davis J, Gallego G, Topaloglu H. 2014. Assortment optimization under variants of the nested logit model. *Operations Research* **62**(2) 250–273.
- Delage E, Kuhn D, Wiesemann W. 2019. “Dice”-sion-making under uncertainty: When can a random decision reduce risk? *Management Science* **65**(7) 3282–3301.
- Delage E, Saif A. 2022. The value of randomized solutions in mixed-integer distributionally robust optimization problems. *INFORMS Journal on Computing* **34**(1) 333–353.
- DeMiguel V, Nogales FJ. 2009. Portfolio selection with robust estimation. *Operations Research* **57**(3) 560–577.
- Désir A, Goyal V, Jiang B, Xie T, Zhang J. 2023. Robust assortment optimization under the Markov chain choice model. *Operations Research* Forthcoming.
- Désir A, Goyal V, Segev D, Ye C. 2020. Constrained assortment optimization under the Markov chain-based choice model. *Management Science* **66**(2) 698–721.
- El Ghaoui L, Lebret H. 1997. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications* **18**(4) 1035–1064.
- Farias V, Jagabathula S, Shah D. 2013. A nonparametric approach to modeling choice with limited data. *Management Science* **59**(2) 305–322.
- Feldman JB, Topaloglu H. 2017. Revenue management under the Markov chain choice model. *Operations Research* **65**(5) 1322–1342.
- Ferreira KJ, Simchi-Levi D, Wang H. 2018. Online network revenue management using Thompson sampling. *Operations Research* **66**(6) 1586–1602.
- Gallego G, Topaloglu H. 2019. *Revenue Management and Pricing Analytics*. Springer.
- Goldfarb D, Iyengar G. 2003. Robust portfolio selection problems. *Mathematics of Operations Research* **28**(1) 1–38.
- Hanasusanto G, Kuhn D, Wiesemann W. 2015. K -adaptability in two-stage robust binary programming. *Operations Research* **63**(4) 877–891.
- Hastie T, Tibshirani R, Friedman J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.

- Honhon D, Jonnalagedda S, Pan XA. 2012. Optimal algorithms for assortment selection under ranking-based consumer choice models. *Manufacturing & Service Operations Management* **14**(2) 279–289.
- Iyengar GN. 2005. Robust dynamic programming. *Mathematics of Operations Research* **30**(2) 257–280.
- Jagabathula S. 2014. Assortment optimization under general choice. Working paper, New York University, New York.
- Jasin S. 2015. Performance of an LP-based control for revenue management with unknown demand parameters. *Operations Research* **63**(4) 909–915.
- Jasin S, Kumar S. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* **37**(2) 313–345.
- Kök AG, Fisher ML, Vaidyanathan R. 2015. Assortment planning: Review of literature and industry practice. Agrawal N, Smith S, eds., *Retail Supply Chain Management: Quantitative Models and Empirical Studies*. Springer, 175–236.
- Li H-L. 1994. A global approach for general 0–1 fractional programming. *European Journal of Operational Research* **73**(3) 590–596.
- Luce RD. 1959. *Individual choice behavior: A theoretical analysis*. Courier Corporation.
- Ma W. 2023. When is assortment optimization optimal? *Management Science* Forthcoming.
- Mahajan S, van Ryzin G. 2001. Stocking retail assortments under dynamic consumer substitution. *Operations Research* **49**(3) 334–351.
- Mak H-Y, Rong Y, Zhang J. 2015. Appointment scheduling with limited distributional information. *Management Science* **61**(2) 316–334.
- Méndez-Díaz I, Miranda-Bront JJ, Vulcano G, Zabala P. 2014. A branch-and-cut algorithm for the latent-class logit assortment problem. *Discrete Applied Mathematics* **164** 246–263.
- Michaud RO. 1989. The Markowitz optimization enigma: Is ‘optimized’ optimal? *Financial Analysts Journal* **45**(1) 31–42.
- Nilim A, El Ghaoui L. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* **53**(5) 780–798.
- Paul A, Feldman J, Davis J. 2018. Assortment optimization and pricing under a nonparametric tree choice model. *Manufacturing & Service Operations Management* **20**(3) 550–565.
- Perakis G, Roels G. 2010. Robust controls for network revenue management. *Manufacturing & Service Operations Management* **12**(1) 56–76.
- Plackett RL. 1975. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **24**(2) 193–202.
- Reiman MI, Wang Q. 2008. An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research* **33**(2) 257–282.

- Rusmevichientong P, Shen Z-JM, Shmoys D. 2010. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* **58**(6) 1666–1680.
- Rusmevichientong P, Shmoys D, Tong C, Topaloglu H. 2014. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* **23**(11) 2023–2039.
- Rusmevichientong P, Topaloglu H. 2012. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research* **60**(4) 865–882.
- Rusmevichientong P, Van Roy B, Glynn PW. 2006. A nonparametric approach to multiproduct pricing. *Operations Research* **54**(1) 82–98.
- Şimşek AS, Topaloglu H. 2018. An expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. *Operations Research* **66**(3) 748–760.
- Smith JE, Winkler RL. 2006. The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. *Management Science* **52**(3) 311–322.
- Talluri K, van Ryzin G. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* **50**(1) 15–33.
- Talluri K, van Ryzin G. 2006. *The Theory and Practice of Revenue Management*. Springer.
- Toubia O, Simester DI, Hauser JR, Dahan E. 2003. Fast polyhedral adaptive conjoint estimation. *Marketing Science* **22**(3) 273–303.
- van Ryzin G, Vulcano G. 2015. A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* **61**(2) 281–300.
- van Ryzin G, Vulcano G. 2017. An expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research* **65**(2) 396–407.
- Wieseemann W, Kuhn D, Rustem B. 2013. Robust Markov decision processes. *Mathematics of Operations Research* **38**(1) 153–183.
- Williams HC. 1977. On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and Planning A* **9**(3) 285–344.
- Xu H, Caramanis C, Mannor S. 2009. Robust regression and Lasso. Koller D, Schuurmans D, Bengio Y, Bottou L, eds., *Advances in Neural Information Processing Systems 21*. Curran Associates, Inc., 1801–1808.
- Zhang D, Cooper WL. 2005. Revenue management for parallel flights with customer-choice behavior. *Operations Research* **53**(3) 415–431.

Appendix A: Exact and Heuristic Solution Schemes for the Randomized Robust Assortment Optimization Problem

We first describe a generic solution scheme for instances $(\mathcal{N}, \mathcal{S}, \mathcal{C}, \mathcal{U}, \mathbf{r})$ of the randomized robust assortment optimization problem **RANDOMIZED ROBUST**. Afterwards, Sections A.1, A.2 and A.3 provide the specific implementation details for the MNL, MC and PR models, respectively.

Column generation scheme for **RANDOMIZED ROBUST**

1. **Initialization.** Set $\hat{\mathcal{S}}$ to any non-empty subset of \mathcal{S} (currently considered assortments). Set $\hat{\mathcal{U}}$ to any non-empty subset of \mathcal{U} (currently considered uncertainty realizations). Set $\text{LB} = -\infty$ and $\text{UB} = +\infty$ (current lower and upper bounds on the optimal value).

2. **Primal Step.** Solve the restricted primal problem

$$\text{maximize } \min_{\mathbf{p} \in \Delta(\hat{\mathcal{S}})} \sum_{\substack{u \in \mathcal{U} \\ S \in \hat{\mathcal{S}}}} p_S \cdot R(S, u) : \quad (11)$$

- (a) Set $\text{LB}_p = -\infty$ and $\text{UB}_p = \text{UB}$ (lower and upper bounds on the optimal value of the restricted primal problem). Set \mathbf{p} to any point in $\Delta(\hat{\mathcal{S}})$ (current restricted randomization strategy).
- (b) Solve the subproblem

$$\tau_{\mathbf{p}}^* = \text{minimize } \sum_{S \in \hat{\mathcal{S}}} p_S \cdot R(S, u). \quad (12)$$

Update $\text{LB}_p = \max\{\text{LB}_p, \tau_{\mathbf{p}}^*\}$ and $\hat{\mathcal{U}} = \hat{\mathcal{U}} \cup \{u^*\}$, where u^* is an optimal solution of the subproblem.

- (c) Solve the doubly restricted primal problem (11) that replaces \mathcal{U} with $\hat{\mathcal{U}}$. Update $\text{UB}_p = \min\{\text{UB}_p, \tau^*\}$, where τ^* is the optimal value of the doubly restricted problem, and set \mathbf{p} to an optimal solution of the problem.
- (d) If $\text{UB}_p - \text{LB}_p < \epsilon/2$, set $\text{LB} = \text{LB}_p$ and go to Step 3. Otherwise, go to Step 2(b).

3. **Dual Step.** Solve the restricted dual problem

$$\text{minimize } \max_{\boldsymbol{\kappa} \in \Delta(\hat{\mathcal{U}})} \sum_{\substack{S \in \mathcal{S} \\ u \in \hat{\mathcal{U}}}} \kappa_u \cdot R(S, u) : \quad (13)$$

- (a) Set $\text{LB}_d = \text{LB}$ and $\text{UB}_d = +\infty$ (lower and upper bounds on the optimal value of the restricted dual problem). Set $\boldsymbol{\kappa}$ to any point in $\Delta(\hat{\mathcal{U}})$ (convex combination of the currently considered valuations).
- (b) Solve the subproblem

$$\tau_{\boldsymbol{\kappa}}^* = \text{maximize } \sum_{\substack{S \in \mathcal{S} \\ u \in \hat{\mathcal{U}}}} \kappa_u \cdot R(S, u). \quad (14)$$

Update $\text{UB}_d = \min\{\text{UB}_d, \tau_{\boldsymbol{\kappa}}^*\}$ and $\hat{\mathcal{S}} = \hat{\mathcal{S}} \cup \{S^*\}$, where S^* is an optimal solution of the subproblem.

- (c) Solve the doubly restricted dual problem (13) that replaces \mathcal{S} with $\hat{\mathcal{S}}$. Update $\text{LB}_d = \max\{\text{LB}_d, \tau^*\}$, where τ^* is the optimal value of the doubly restricted problem, and set $\boldsymbol{\kappa}$ to an optimal solution of the problem.
- (d) If $\text{UB}_d - \text{LB}_d < \epsilon/2$, set $\text{UB} = \text{UB}_d$ and go to Step 4. Otherwise, go to Step 3(b).

4. **Termination.** Iterate between Steps 2 and 3 until the gap between LB and UB is zero or sufficiently small. Then terminate with the primal-dual pair $(\mathbf{p}^*, \boldsymbol{\kappa}^*) \in \mathbb{R}^{|\mathcal{S}|} \times \mathbb{R}^{|\mathcal{U}|}$, where $p_S^* = p_S$, $S \in \hat{\mathcal{S}}$, $= 0$ otherwise and $\kappa_u^* = \kappa_u$, $u \in \hat{\mathcal{U}}$, $= 0$ otherwise.

The ‘outer layer’ of the algorithm iterates between conservative approximations to the primal and dual problems (11) and (13) that consider all uncertainty realizations $u \in \mathcal{U}$ (in the primal problem) and assortments $S \in \mathcal{S}$ (in the dual problem) but that only consider subsets of the admissible decisions p_S , $S \in \hat{\mathcal{S}} \subseteq \mathcal{S}$ (in the primal problem) and κ_u , $u \in \hat{\mathcal{U}} \subseteq \mathcal{U}$ (in the dual problem). The resulting restrictions remain computationally challenging as they involve embedded non-convex optimization problems over $u \in \mathcal{U}$ and $S \in \mathcal{S}$, respectively, and they are solved iteratively up to a pre-specified tolerance of $\epsilon \geq 0$. To this end, the ‘inner layers’ of the algorithm iterate between (i) determining an optimal solution to the embedded optimization problem while fixing the solution to the outer optimization problems in Steps 2(b) and 3(b) and (ii) solving the outer optimization problems for the refined approximations of the embedded optimization problems in Steps 2(c) and 3(c).

The correctness of the algorithm is summarized in the following statement.

THEOREM 8. *The two-layer column generation algorithm presented above converges in a finite number of iterations and terminates with an optimal primal-dual solution pair $(\mathbf{p}^*, \boldsymbol{\kappa}^*)$ to problems (11) and (13).*

Proof of Theorem 8. The proof largely follows similar arguments as the proof of Theorem 2 of Delage and Saif (2022) and is thus omitted for the sake of brevity. \square

Our column generation scheme is reminiscent to, but differs in several important aspects from, the two-layer algorithm proposed in Section 5.3 of Delage and Saif (2022). While Delage and Saif (2022) study a *two-stage* robust optimization problem with a *linear* objective function, our randomized assortment optimization problem constitutes a *single-stage* robust optimization problem that comprises a *non-convex* objective function. To solve their problem, Delage and Saif (2022) equivalently express their primal and dual subproblems through Karush-Kuhn-Tucker reformulations, whereas our subproblems require a different case-by-case treatment for each choice model that we elaborate upon in the following subsections.

A.1. Exact and Heuristic Solution of the Randomized Constrained MNL Problem

Recall from equation (2) that for the MNL choice model, the revenues $R(S, u)$ evaluate to

$$R(S, u) = R(S, \mathbf{v}) = \sum_{i \in S} r_i \cdot \psi_i(S, \mathbf{v}) = \frac{\sum_{i \in S} r_i \cdot v_i}{v_0 + \sum_{i \in S} v_i},$$

and recall from equation (3) that we consider the binary representable uncertainty set

$$\mathcal{V} = \{\mathbf{v} = \mathbf{F}\boldsymbol{\xi} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}, \boldsymbol{\xi} \in \{0, 1\}^m\}.$$

In the following, we present mixed-integer linear programs (MILPs) for the subproblems (12) and (14) of our column generation scheme that resemble the MILP formulations for sum-of-ratio problems, see Li (1994).

For the subproblem (12) in the primal step of our column generation scheme, we first replace the decision variables $\mathbf{v} \in \mathcal{V}$ with their definition and introduce auxiliary variables d_S , $S \in \hat{\mathcal{S}}$, that equate to the (inverses of the) denominators in the objective function:

$$\begin{aligned} & \text{minimize} && \sum_{S \in \hat{\mathcal{S}}} p_S \cdot d_S \cdot \sum_{i \in S} r_i \cdot \mathbf{F}_i^\top \boldsymbol{\xi} \\ & \text{subject to} && d_S \cdot \left[\sum_{i \in S \cup \{0\}} \mathbf{F}_i^\top \boldsymbol{\xi} \right] = 1 && \forall S \in \hat{\mathcal{S}} \\ & && \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b} \\ & && \boldsymbol{\xi} \in \{0, 1\}^m, d_S \in \mathbb{R}_+, S \in \hat{\mathcal{S}}. \end{aligned}$$

Here, \mathbf{F}_i^\top denotes the i -th row of the matrix \mathbf{F} (where the counting starts from zero). For all $S \in \hat{\mathcal{S}}$ and $j \in \{1, \dots, m\}$, we now apply the exact linearization

$$z_{Sj} = d_S \cdot \xi_j \iff \begin{bmatrix} z_{Sj} \leq d_S, & z_{Sj} \leq M \cdot \xi_j \\ z_{Sj} \geq 0, & z_{Sj} \geq d_S - M \cdot (1 - \xi_j) \end{bmatrix},$$

where $z_{Sj} \in \mathbb{R}$, $S \in \hat{\mathcal{S}}$ and $j = 1, \dots, m$, constitute auxiliary decision variables and M denotes a sufficiently large positive constant that can be determined from the geometry of the uncertainty set (Li 1994). The resulting formulation is an MILP that can be solved with off-the-shelf software.

The subproblem (14) in the dual step of our algorithm constitutes a nominal assortment optimization problem under the mixture of multinomial logits (MMNL) choice model. MILP reformulations for this problem have been proposed, among others, by Bront et al. (2009), Méndez-Díaz et al. (2014) and Şen et al. (2018). To solve this problem, we first introduce the binary decision variables $x_i \in \{0, 1\}$, $i \in \mathcal{N}$, to denote whether product i is contained in the assortment S and introduce auxiliary variables d_v , $\mathbf{v} \in \hat{\mathcal{V}}$, that equate to the (inverses of the) denominators in the objective function:

$$\begin{aligned} & \text{maximize} && \sum_{\mathbf{v} \in \hat{\mathcal{V}}} \kappa_{\mathbf{v}} \cdot d_{\mathbf{v}} \cdot \sum_{i \in \mathcal{N}} r_i \cdot v_i \cdot x_i \\ & \text{subject to} && d_{\mathbf{v}} \cdot \left[v_0 + \sum_{i \in \mathcal{N}} v_i \cdot x_i \right] = 1 \quad \forall \mathbf{v} \in \hat{\mathcal{V}} \\ & && \sum_{i \in \mathcal{N}} x_i \leq C \\ & && \mathbf{x} \in \{0, 1\}^n, d_{\mathbf{v}} \in \mathbb{R}_+, \mathbf{v} \in \hat{\mathcal{V}} \end{aligned}$$

For all $\mathbf{v} \in \hat{\mathcal{V}}$ and $i \in \mathcal{N}$, we now apply the exact linearization

$$z_{vi} = d_{\mathbf{v}} \cdot x_i \iff \begin{bmatrix} z_{vi} \leq d_{\mathbf{v}}, & z_{vi} \leq M \cdot x_i \\ z_{vi} \geq 0, & z_{vi} \geq d_{\mathbf{v}} - M \cdot (1 - x_i) \end{bmatrix},$$

where $z_{vi} \in \mathbb{R}$, $\mathbf{v} \in \hat{\mathcal{V}}$ and $i \in \mathcal{N}$, constitute auxiliary decision variables and M denotes a sufficiently large positive constant that can again be determined from the geometry of the uncertainty set (Li 1994). The resulting formulation is again an MILP.

The main computational challenge in the column generation scheme of the previous section arises from the subproblems of Steps 2(b) and 3(b), which amount to discrete sums-of-ratios problems. As an alternative to solving those problems exactly, we can also adapt the ADXOpt heuristic of Jagabathula (2014), which was originally designed for nominal assortment optimization problems under general choice models, to these subproblems. Since the subproblem of Step 3(b) is already in the form considered by Jagabathula (2014), we focus on the subproblem of Step 2(b) and consider the budget uncertainty set of Section 4.2:

$$\mathcal{V} = \{ \mathbf{v} = \bar{\mathbf{v}} - (\bar{\mathbf{v}} - \underline{\mathbf{v}}) \circ \boldsymbol{\xi} : \mathbf{e}^\top \boldsymbol{\xi} \leq \Gamma, \boldsymbol{\xi} \in \{0, 1\}^{n+1} \}$$

The adaptations to other uncertainty sets are straightforward.

Local search heuristic for problem (12)

1. **Initialization.** Set $\boldsymbol{\xi} = \mathbf{0}$ (current worst-case valuation).

2. **Additions.** If $\mathbf{e}^\top \boldsymbol{\xi} < \Gamma$, evaluate the objective function of problem (12) for all $\boldsymbol{\xi}' = \boldsymbol{\xi} + \mathbf{e}_i$, $i \in \mathcal{N}_0$, for which $\boldsymbol{\xi}' \in \{0, 1\}^{n+1}$. If any such $\boldsymbol{\xi}'$ improves the objective function of (12), update $\boldsymbol{\xi}$ to the best such $\boldsymbol{\xi}'$ and repeat Step 2.
3. **Deletions and exchanges.** Evaluate the objective function of problem (12) for all $\boldsymbol{\xi}' = \boldsymbol{\xi} - \mathbf{e}_i$, $i \in \mathcal{N}_0$, and $\boldsymbol{\xi}' = \boldsymbol{\xi} + \mathbf{e}_i - \mathbf{e}_j$, $i, j \in \mathcal{N}_0$ and $i \neq j$, for which $\boldsymbol{\xi}' \in \{0, 1\}^{n+1}$. If any such $\boldsymbol{\xi}'$ improves the objective function of (12), update $\boldsymbol{\xi}$ to the best such $\boldsymbol{\xi}'$ and go to Step 2. Otherwise, go to Step 4.
4. **Termination.** Terminate and return $\boldsymbol{\xi}$ as worst-case valuation.

Being a local search heuristic, the worst-case valuation returned by our adaptation of ADXOpt may not represent a global minimizer of the primal subproblem (12). As such, we may increase LB_p and, subsequently, LB , excessively. Likewise, the application of ADXOpt to the dual subproblem (14) may decrease UB_d and, subsequently, UB , excessively. Either of these occurrences may cause the two-layer column generation scheme to terminate early with a feasible but suboptimal primal-dual solution pair $(\mathbf{p}^*, \boldsymbol{\kappa}^*)$. Our numerical experiments in Section 7 show, however, that the local search heuristic performs remarkably well and returns solutions that are very close—and often identical—to the optimal ones.

A.2. Exact and Heuristic Solution of the Randomized MC Problem

For the MC choice model, the revenues $R(S, u)$ can be expressed as

$$R(S, u) = R(S, \boldsymbol{\rho}) = \sum_{i \in \mathcal{N}} \lambda_i \cdot g_i \quad \text{with} \quad g_i = \begin{cases} r_i & \text{if } i \in S, \\ \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j & \text{if } i \in \mathcal{N} \setminus S. \end{cases}$$

Also, recall that we consider finite uncertainty sets $\mathcal{U} \subseteq \mathbb{R}_+^{(n+1) \times (n+1)}$ for the general substitution sets as well as $\mathcal{U} = \prod_{i \in \mathcal{N}_0} \mathcal{U}_i$ with $\mathcal{U}_i \subseteq \mathbb{R}^{n+1}$, $i \in \mathcal{N}_0$, for the product-wise substitution sets.

To solve the subproblem (12) in the primal step of our column generation scheme, we first consider the case of general substitution sets. In that case, we can explicitly evaluate $R(S, \boldsymbol{\rho})$ for each assortment $S \in \hat{\mathcal{S}}$ and each substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$ by solving a separate linear program

$$\begin{aligned} & \text{minimize} && \sum_{i \in \mathcal{N}} \lambda_i \cdot g_i \\ & \text{subject to} && g_i = r_i && \forall i \in S \\ & && g_i \geq \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j && \forall i \in \mathcal{N} \setminus S \\ & && g_i \in \mathbb{R}, i \in \mathcal{N} \end{aligned}$$

and subsequently select a substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$ that minimizes $\sum_{S \in \hat{\mathcal{S}}} p_S \cdot R(S, \boldsymbol{\rho})$. A similar approach would be inefficient for product-wise substitution sets, however, as $\prod_{i \in \mathcal{N}_0} |\mathcal{U}_i|$ many linear programs would have to be solved. Instead, we first formulate the subproblem (12) as the mixed-integer bilinear program

$$\begin{aligned}
& \text{minimize} && \sum_{S \in \hat{\mathcal{S}}} p_S \cdot \left[\sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^S \right] \\
& \text{subject to} && g_i^S = r_i && \forall i \in S, \forall S \in \hat{\mathcal{S}} \\
& && g_i^S \geq \sum_{\rho_i \in \mathcal{U}_i} x_i^{\rho_i} \cdot y_i^{\rho_i S} && \forall i \in \mathcal{N} \setminus S, \forall S \in \hat{\mathcal{S}} \\
& && y_i^{\rho_i S} = \sum_{j \in \mathcal{N}} \rho_{ij}^{\rho_i} \cdot g_j^S && \forall i \in \mathcal{N} \setminus S, \forall S \in \hat{\mathcal{S}}, \forall \rho_i \in \mathcal{U}_i \\
& && \sum_{\rho_i \in \mathcal{U}_i} x_i^{\rho_i} = 1 && \forall i \in \mathcal{N} \\
& && g_i^S \in \mathbb{R}, i \in \mathcal{N} \text{ and } S \in \hat{\mathcal{S}} \\
& && x_i^{\rho_i} \in \{0, 1\}, i \in \mathcal{N} \text{ and } \rho_i \in \mathcal{U}_i \\
& && y_i^{\rho_i S} \in \mathbb{R}, i \in \mathcal{N}, \rho_i \in \mathcal{U}_i \text{ and } S \in \hat{\mathcal{S}},
\end{aligned}$$

where $x_i^{\rho_i}$ is 1 if $\rho_i \in \mathcal{U}_i$ is selected as worst-case substitution vector for product $i \in \mathcal{N}$, and 0 otherwise. Note that the variables $x_i^{\rho_i}$ satisfying $i \in \bigcap_{S \in \hat{\mathcal{S}}} S$ do not appear in the first three constraints and can thus be eliminated from the problem. For all $\rho_i \in \mathcal{U}_i$, $i \in \mathcal{N} \setminus S$ and $S \in \hat{\mathcal{S}}$, we now apply the exact linearization

$$z_i^{\rho_i S} = x_i^{\rho_i} \cdot y_i^{\rho_i S} \iff \begin{bmatrix} z_i^{\rho_i S} \leq x_i^{\rho_i}, & z_i^{\rho_i S} \leq y_i^{\rho_i S} \\ z_i^{\rho_i S} \geq 0, & z_i^{\rho_i S} \geq x_i^{\rho_i} + y_i^{\rho_i S} - 1 \end{bmatrix},$$

where $z_i^{\rho_i S} \in \mathbb{R}$, $i \in \mathcal{N} \setminus S$, $S \in \hat{\mathcal{S}}$ and $\rho_i \in \mathcal{U}_i$, constitute auxiliary decision variables. The resulting formulation is an MILP that can be solved with off-the-shelf software.

In view of the subproblem (14) in the dual step of our algorithm, we introduce the binary decision variables $x_i \in \{0, 1\}$, $i \in \mathcal{N}$, to denote whether product i is contained in the assortment S , as well as the continuous auxiliary decision variables g_i^ρ , $i \in \mathcal{N}$ and $\rho \in \hat{\mathcal{U}}$, to obtain the following mixed-integer bilinear program:

$$\begin{aligned}
& \text{maximize} && \sum_{\rho \in \hat{\mathcal{U}}} \kappa_\rho \cdot \left[\sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^\rho \right] \\
& \text{subject to} && g_i^\rho \leq r_i \cdot x_i + \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j^\rho \cdot (1 - x_i) && \forall i \in \mathcal{N}, \rho \in \hat{\mathcal{U}} \\
& && g_i^\rho \in \mathbb{R}, i \in \mathcal{N} \text{ and } \rho \in \hat{\mathcal{U}} \\
& && \mathbf{x} \in \mathcal{X}
\end{aligned}$$

Here, $\mathcal{X} \subseteq \{0, 1\}^n$ reflects the business constraints encoded by the choice of \mathcal{S} . The first constraint enforces that under the substitution matrix $\rho \in \hat{\mathcal{U}}$, g_i^ρ attains the value r_i if product i is selected, and g_i^ρ is upper bounded by $\sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j^\rho$ otherwise. We apply the exact linearization

$$z_i^\rho = x_i \cdot y_i^\rho \iff \begin{bmatrix} z_i^\rho \leq x_i, & z_i^\rho \leq y_i^\rho \\ z_i^\rho \geq 0, & z_i^\rho \geq x_i + y_i^\rho - 1 \end{bmatrix},$$

where $z_i^\rho \in \mathbb{R}$, $i \in \mathcal{N}$ and $\rho \in \hat{\mathcal{U}}$, are auxiliary decision variables, to obtain an equivalent MILP reformulation.

As in the MNL choice model, the main computational challenge in the column generation scheme arises from the subproblems of Steps 2(b) and 3(b). Instead of solving those problem exactly, we can principally solve them with local search heuristics akin to the ADXOpt heuristic of Jagabathula (2014). In our numerical experiments, however, such local search heuristics did not perform well as the evaluation of each candidate solution requires the solution of a linear program for each $S \in \hat{\mathcal{S}}$ (for the primal subproblem, to determine the values of g_i^S) or $\rho \in \hat{\mathcal{U}}$ (for the dual subproblem, to determine the values of g_i^ρ). This is in contrast to the MNL and PR models, where the evaluation of each candidate solution only requires the computation of a closed-form expression.

A.3. Exact and Heuristic Solution of the Randomized PR Problem

Recall from equation (7) that for the PR model, the revenues $R(S, u)$ evaluate to

$$R(S, u) = R(S, \boldsymbol{\lambda}) = \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S),$$

and recall from equation (8) that we consider the polyhedral uncertainty set

$$\Lambda = \{ \boldsymbol{\lambda} = \mathbf{F}\boldsymbol{\xi} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}, \boldsymbol{\xi} \in \mathbb{R}_+^m \}.$$

The randomized assortment optimization problem under the PR model is simpler than its counterparts under the MNL and MC models since in the present case, the revenues are linear in the uncertainty realization. In the following, we provide a simplified column generation scheme that is reminiscent of the single-layer algorithm proposed in Section 5.4 of [Delage and Saif \(2022\)](#).

Column generation scheme for problem (9)

1. **Initialization.** Set $\hat{\mathcal{S}}$ to any non-empty subset of \mathcal{S} (currently considered assortments). Set $\text{LB} = -\infty$ and $\text{UB} = +\infty$ (current lower and upper bounds on the optimal value).
2. **Master problem.** Solve the restricted dual problem (10) that replaces $S \in \mathcal{S}$ with $S \in \hat{\mathcal{S}}$ in the objective function. Update $\text{LB} = \max\{\text{LB}, \tau^*\}$, where τ^* is the optimal value of the restricted dual problem, and set $\boldsymbol{\lambda}$ to an optimal solution of the problem.
3. **Subproblem.** Solve the subproblem

$$\begin{aligned} & \text{maximize} && \sum_{k \in \mathcal{K}} \lambda_k \cdot R_k(S) \\ & \text{subject to} && S \in \mathcal{S}. \end{aligned}$$

Update $\text{UB} = \min\{\text{UB}, \tau^*\}$, where τ^* is the optimal value of the subproblem, and $\hat{\mathcal{S}} = \hat{\mathcal{S}} \cup \{S^*\}$, where S^* is an optimal solution of the problem.

4. **Termination.** Iterate between Steps 2 and 3 until the gap between LB and UB is zero or sufficiently small. Then terminate with the primal-dual pair $(\mathbf{p}^*, \boldsymbol{\lambda}^*) \in \mathbb{R}^{|\mathcal{S}|} \times \mathbb{R}^K$, where p_S^* , $S \in \hat{\mathcal{S}}$, can be derived from the shadow prices of the epigraph reformulation of the restricted dual problem (10), and $p_S^* = 0$ for $S \notin \hat{\mathcal{S}}$.

The algorithm iterates between solving two problems: Step 2 solves an increasingly accurate progressive approximation of the dual problem (10) to determine a candidate solution $\boldsymbol{\lambda} \in \Lambda$. Step 3 evaluates the exact objective function of problem (10) for $\boldsymbol{\lambda}$ and identifies a cutting plane to add to the approximation from Step 2. Both problems can be expressed as finite (MI)LPs that can be solved with standard software ([Bertsimas and Mišić 2019](#)).

The correctness of the algorithm is summarized next.

THEOREM 9. *The column generation algorithm presented above converges in a finite number of iterations and terminates with an optimal primal-dual solution pair $(\mathbf{p}^*, \boldsymbol{\lambda}^*)$ to problems (9) and (10).*

Proof of Theorem 9. The proof largely follows similar arguments as the proof of Theorem 2 of [Delage and Saif \(2022\)](#) and is thus omitted for the sake of brevity. \square

The main computational challenge in the above column generation scheme arises from the subproblem of Step 3, which amounts to a nominal assortment optimization problem over the PR model. Instead of solving this problem exactly, we can solve it heuristically with the ADXOpt heuristic of [Jagabathula \(2014\)](#). In doing so, we may miss the global maximizer of the subproblem and therefore decrease UB excessively, which in turn can cause the column generation scheme to terminate early with a feasible but suboptimal primal-dual solution pair $(\mathbf{p}^*, \boldsymbol{\lambda}^*)$. Similar to the cardinality-constrained MNL model, however, our numerical experiments in [Section 7](#) show that the local search heuristic performs remarkably well and generally returns solutions of excellent quality.

Appendix B: Proofs for Section 3

Proof of Theorem 1. Denote by $\mathcal{R} = \{\mathbf{R} \in \mathbb{R}^{|\mathcal{S}|} : [R_S = R(S, u) \forall S \in \mathcal{S}] \text{ for some } u \in \mathcal{U}\}$ the set of all expected revenues vectors for the different uncertainty realizations $u \in \mathcal{U}$. The statement of the theorem then claims that the instance $(\mathcal{N}, \mathcal{S}, \mathcal{C}, \mathcal{U}, \mathbf{r})$ is randomization-proof if and only if there is $\mathbf{R}' \in \text{conv } \mathcal{R}$ such that $\mathbf{R}' \leq R_{\text{det}}^*(\mathcal{U}) \cdot \mathbf{e}$, where $\text{conv } \mathcal{R}$ denotes the convex hull of \mathcal{R} .

We first note that

$$R^*(\mathbf{p}) = \min_{u \in \mathcal{U}} R(\mathbf{p}, u) = \min_{\mathbf{R} \in \mathcal{R}} \mathbf{p}^\top \mathbf{R} = \min_{\mathbf{R} \in \text{conv } \mathcal{R}} \mathbf{p}^\top \mathbf{R}, \quad (15)$$

where the last identity follows from the fact that the minimization over $\mathbf{R} \in \mathcal{R}$ has a linear objective, which allows us to replace its feasible region \mathcal{R} by its convex hull $\text{conv } \mathcal{R}$.

Assume first that there is $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R_{\text{det}}^*(\mathcal{U})$ for all $S \in \mathcal{S}$. As discussed before, this is equivalent to assuming that there is $\mathbf{R}' \in \text{conv } \mathcal{R}$ such that $\mathbf{R}' \leq R_{\text{det}}^*(\mathcal{U}) \cdot \mathbf{e}$. In that case, any randomized assortment $\mathbf{p} \in \Delta(\mathcal{S})$ satisfies

$$R^*(\mathbf{p}) = \min_{\mathbf{R} \in \text{conv } \mathcal{R}} \mathbf{p}^\top \mathbf{R} \leq \mathbf{p}^\top \mathbf{R}' \leq R_{\text{det}}^*(\mathcal{U}),$$

where the identity follows from (15), the first inequality holds since $\mathbf{R}' \in \text{conv } \mathcal{R}$, and the second inequality is due to the fact that $\mathbf{p} \in \Delta(\mathcal{S})$ and $\mathbf{R}' \leq R_{\text{det}}^*(\mathcal{U}) \cdot \mathbf{e}$.

Assume now that there is no $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R_{\text{det}}^*(\mathcal{U})$ for all $S \in \mathcal{S}$. This is equivalent to assuming that there is no $\mathbf{R}' \in \text{conv } \mathcal{R}$ such that $\mathbf{R}' \leq R_{\text{det}}^*(\mathcal{U}) \cdot \mathbf{e}$. In that case, we note that

$$\max_{\mathbf{p} \in \Delta(\mathcal{S})} R^*(\mathbf{p}) = \max_{\mathbf{p} \in \Delta(\mathcal{S})} \min_{\mathbf{R} \in \text{conv } \mathcal{R}} \mathbf{p}^\top \mathbf{R} = \min_{\mathbf{R} \in \text{conv } \mathcal{R}} \max_{\mathbf{p} \in \Delta(\mathcal{S})} \mathbf{p}^\top \mathbf{R} = \min_{\mathbf{R} \in \text{conv } \mathcal{R}} \max_{S \in \mathcal{S}} R_S > R_{\text{det}}^*(\mathcal{U}),$$

where the first identity follows from (15), the second one is due to strong duality, the third one holds since the objective function of the inner maximization is linear in \mathbf{p} , and the inequality holds by assumption. \square

OBSERVATION 6. *The condition in Theorem 1 is equivalent to requiring that at every optimal deterministic assortment $S^* \in \mathcal{S}$, the super-differential of the graph $\mathbf{p} \mapsto R^*(\mathbf{p})$ at $\mathbf{p} = \mathbf{e}_{S^*}$ contains a supergradient \mathbf{R} such that $\mathbf{R}^\top (\mathbf{p} - \mathbf{e}_{S^*}) \leq 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$.*

Since the graph $\mathbf{p} \mapsto R^*(\mathbf{p})$ is concave over $\Delta(\mathcal{S})$, Observation 6 expresses our intuition from the main paper that the graph $\mathbf{p} \mapsto R^*(\mathbf{p})$ does not contain a feasible direction of improvement at $\mathbf{p} = \mathbf{e}_{S^*}$.

Proof of Observation 6. Fix any optimal deterministic assortment $S^* \in \mathcal{S}$. Using the definition of $R^*(\mathbf{p})$, one can verify that the super-differential of R^* satisfies

$$\partial R^*(\mathbf{e}_{S^*}) = \text{conv} \{ \mathbf{R} \in \mathbb{R}^{|\mathcal{S}|} : [R_S = R(S, u) \forall S \in \mathcal{S}] \text{ for } u \in \mathcal{U} \text{ such that } R(S^*, u) = R^*(\mathbf{e}_{S^*}) = R_{\text{det}}^*(\mathcal{U}) \}.$$

We show in three steps that there is $\mathbf{R} \in \partial R^*(\mathbf{e}_{S^*})$ such that $\mathbf{R}^\top (\mathbf{p} - \mathbf{e}_{S^*}) \leq 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$ if and only if there is $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R^*(S^*)$ for all $S \in \mathcal{S}$. In Step 1, we argue that for a suitably defined set \mathcal{U}^* there is a one-to-one correspondence between the convex combinations generating $\mathbf{R} \in \partial R^*(\mathbf{e}_{S^*})$ and the probability distributions $\mathbb{Q} \in \mathcal{P}(\mathcal{U}^*)$ satisfying $R_S = \mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})]$ for all $S \in \mathcal{S}$. In Step 2, we show that \mathbf{R} satisfies $\mathbf{R}^\top (\mathbf{p} - \mathbf{e}_{S^*}) \leq 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$ if and only if $R_S \leq R^*(S^*)$ for all $S \in \mathcal{S}$. Taken together, Steps 1

and 2 imply that there is $\mathbf{R} \in \partial R^*(\mathbf{e}_{S^*})$ such that $\mathbf{R}^\top(\mathbf{p} - \mathbf{e}_{S^*}) \leq 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$ if and only if there is $\mathbb{Q} \in \mathcal{P}(\mathcal{U}^*)$ such that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R^*(S^*)$ for all $S \in \mathcal{S}$. Step 3, finally, shows that this conclusion continues to hold if we replace the set \mathcal{U}^* with the set \mathcal{U} .

In view of Step 1, let \mathcal{U}^* denote the set of all uncertainty realizations $u \in \mathcal{U}$ for which $R(S^*, u) = R_{\det}^*(\mathcal{U})$. One then verifies that there is a one-to-one correspondence between the convex combinations generating the vectors $\mathbf{R} \in \partial R^*(\mathbf{e}_{S^*})$ and the probability distributions $\mathbb{Q} \in \mathcal{P}(\mathcal{U}^*)$ satisfying $R_S = \mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})]$ for all $S \in \mathcal{S}$.

As for Step 2, note that \mathbf{R} satisfies $\mathbf{R}^\top(\mathbf{p} - \mathbf{e}_{S^*}) \leq 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$ if and only if $R_S \leq R_{\det}^*(\mathcal{U})$ for all $S \in \mathcal{S}$. Indeed, the ‘only if’ direction follows from the choices $\mathbf{p} = \mathbf{e}_S$, $S \in \mathcal{S}$, and the ‘if’ direction holds since $\mathbf{R}^\top(\mathbf{p} - \mathbf{e}_{S^*}) \leq [R_{\det}^*(\mathcal{U}) \cdot \mathbf{e}]^\top(\mathbf{p} - \mathbf{e}_{S^*}) = 0$ for all $\mathbf{p} \in \Delta(\mathcal{S})$.

In view of Step 3, finally, note that $\mathcal{U}^* \subseteq \mathcal{U}$, and thus the statement holds as long as every $\mathbb{Q} \in \mathcal{P}(\mathcal{U})$ satisfying $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R^*(S^*)$ for all $S \in \mathcal{S}$ also satisfies $\mathbb{Q}[\tilde{u} \in \mathcal{U} \setminus \mathcal{U}^*] = 0$. Indeed, if $\mathbb{Q}[\tilde{u} \in \mathcal{U} \setminus \mathcal{U}^*] > 0$, then $\mathbb{E}_{\mathbb{Q}}[R(S^*, \tilde{u})] > R^*(S^*)$, which contradicts the assumption that $\mathbb{E}_{\mathbb{Q}}[R(S, \tilde{u})] \leq R^*(S^*)$ for all $S \in \mathcal{S}$. \square

Proof of Corollary 1. Fix $u' \in \arg \min_{u \in \mathcal{U}} \max_{S \in \mathcal{S}} R(S, u)$ and $S' \in \arg \max_{S \in \mathcal{S}} R(S, u')$. For any $\mathbf{p} \in \Delta(\mathcal{S})$ in (RANDOMIZED ROBUST), we then have

$$R^*(\mathbf{p}) = \min_{u \in \mathcal{U}} \sum_{S \in \mathcal{S}} p_S \cdot R(S, u) \leq \sum_{S \in \mathcal{S}} p_S \cdot R(S, u') \leq \sum_{S \in \mathcal{S}} p_S \cdot R(S', u') = R(S', u') = \max_{S \in \mathcal{S}} \min_{u \in \mathcal{U}} R(S, u),$$

where the first inequality holds since $u' \in \mathcal{U}$ and the second inequality holds by construction of S' , respectively. The second identity follows from the fact that $\sum_{S \in \mathcal{S}} p_S = 1$. The last identity, finally, follows from the construction of S' and u' as well as the assumed strong duality. The statement now follows since the choice of the probability vector \mathbf{p} was arbitrary. \square

Proof of Corollary 2. Assume that the worst-case parameter set is indeed a singleton. If the inequality from the statement of the corollary holds, then the choice $\mathbb{Q} = \delta_{u^*}$ in Theorem 1, which is the Dirac distribution that places unit mass on the realization $\tilde{u} = u^*$, implies that the problem is randomization-proof. If the problem is randomization-proof, on the other hand, then any distribution \mathbb{Q} satisfying the condition in Theorem 1 has to satisfy $\mathbb{Q}[\tilde{u} = u^*] = 1$ since $R(S^*, u) > R^*(S^*)$ for all $u \in \mathcal{U} \setminus \{u^*\}$. In that case, however, we have $R(S, u^*) \leq R^*(S^*)$ for all $S \in \mathcal{S}$, that is, the inequality from the statement of the corollary holds. \square

Appendix C: Proofs for Section 4

For our analysis of the cardinality-constrained MNL model, we consider for any number of products $n \geq 2$ and any cardinality restriction $|S| \leq C$, $C \in \{1, \dots, n-1\}$ the instance where the price of each product i , $i \in \{1, \dots, n\}$, is $r_i = 1$, while the uncertainty set is $\mathcal{V} = \{(1, \mathbf{v}) : \sum_{i \in \mathcal{N}} v_i = n - C, \mathbf{v} \in \{0, 1\}^n\}$. In the remainder of this section, we refer to this model as the *reference model*.

The proof of [Theorem 2](#) relies on the following auxiliary results, which we state and prove first.

LEMMA 1. *Under the reference model, the worst-case expected revenues of any feasible solution to the deterministic robust assortment optimization problem are 0.*

Proof of Lemma 1. Fix any feasible solution $S \in \mathcal{S}$, $|S| \leq C$, to the deterministic robust assortment optimization problem. Under any valuation scenario $\mathbf{v} = (1, \mathbf{v}') \in \mathcal{V}$ with $v'_i = 0$ if $i \in S$, the expected revenues of S vanish. The statement now follows from the fact that the worst-case expected revenues are always non-negative. \square

LEMMA 2. *Under the reference model, the randomized assortment strategy $\mathbf{p}^* \in \Delta(\mathcal{S})$, where $p_S^* = 1/|\{S \in \mathcal{S} : |S| = C\}|$ if $|S| = C$ and $p_S^* = 0$ otherwise, strictly outperforms any other feasible randomized assortment strategy \mathbf{p} that only places positive probability $p_S > 0$ on assortments S of full cardinality $|S| = C$.*

Proof of Lemma 2. Assume to the contrary that there is a feasible randomized assortment strategy $\mathbf{p} \neq \mathbf{p}^*$ that (i) places positive probability $p_S > 0$ only on assortments S of cardinality $|S| = C$ and that (ii) has worst-case revenues larger than or equal to those of \mathbf{p}^* . Since $\mathbf{p} \neq \mathbf{p}^*$, there must be two components i and j of \mathbf{p} —without loss of generality $i = 1$ and $j = 2$ —such that $p_i > p_j$ and the associated assortments S_i and S_j satisfy $|S_i| = |S_j| = C$. Let Π denote the group of all permutations $\pi : \mathbb{R}^{2^n} \rightarrow \mathbb{R}^{2^n}$ that permute those components of a randomization strategy in $\Delta(\mathcal{S})$ that correspond to assortments S of cardinality $|S| = C$. Moreover, assume that $\pi_1 \in \Pi$ is the identity and that $\pi_2 \in \Pi$ reverses the order of the first two components but leaves all other components unchanged.

Consider now the randomization strategy \mathbf{p}' defined through

$$\begin{aligned} \mathbf{p}' &= \frac{1}{|\Pi \setminus \{\pi_1, \pi_2\}|} \sum_{\pi \in \Pi \setminus \{\pi_1, \pi_2\}} \pi(\mathbf{p}) + \left(\frac{1}{|\Pi|} + \frac{\epsilon}{p_1 - p_2} \right) \cdot \pi_1(\mathbf{p}) + \left(\frac{1}{|\Pi|} - \frac{\epsilon}{p_1 - p_2} \right) \cdot \pi_2(\mathbf{p}) \\ &= \mathbf{p}^* + \epsilon(+1, -1, 0, \dots, 0)^\top. \end{aligned}$$

For $\epsilon > 0$ sufficiently small, \mathbf{p}' is a convex combination of the permutations $\pi(\mathbf{p})$, $\pi \in \Pi$, and as such $\mathbf{p}' \in \Delta(\mathcal{S})$. Moreover, we have $p'_S > 0$ only for assortments $S \in \mathcal{S}$ with $|S| = C$. By the symmetry of the problem instance, any permutation $\pi(\mathbf{p})$, $\pi \in \Pi$, has the same worst-case expected revenues as \mathbf{p} . The concavity of the worst-case expected revenues then implies that \mathbf{p}' has worst-case expected revenues that are at least as large as those of \mathbf{p} and thus—by our earlier assumption—of \mathbf{p}^* . We next show that this leads to a contradiction, which will complete the proof.

To compare the worst-case expected revenues of \mathbf{p}' and \mathbf{p}^* , we note that

$$\min_{\mathbf{v} \in \mathcal{V}} \sum_{S \in \mathcal{S}} p'_S \cdot R(S, \mathbf{v}) - \min_{\mathbf{v} \in \mathcal{V}} \sum_{S \in \mathcal{S}} p_S^* \cdot R(S, \mathbf{v})$$

$$\begin{aligned}
&= \min_{\mathbf{v} \in \mathcal{V}} \sum_{S \in \mathcal{S}} [p'_S - p_S^*] \cdot R(S, \mathbf{v}) \\
&= \min_{\mathbf{v} \in \mathcal{V}} \epsilon \cdot [R(S_1, \mathbf{v}) - R(S_2, \mathbf{v})] \\
&= \min_{\mathbf{v} \in \mathcal{V}} \epsilon \cdot \left[\frac{\sum_{i \in S_1} v_i}{v_0 + \sum_{i \in S_1} v_i} - \frac{\sum_{i \in S_2} v_i}{v_0 + \sum_{i \in S_2} v_i} \right] \\
&= \min_{\mathbf{v} \in \mathcal{V}} \epsilon \cdot \left[\frac{\sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_1 \setminus S_2} v_i}{v_0 + \sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_1 \setminus S_2} v_i} - \frac{\sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_2 \setminus S_1} v_i}{v_0 + \sum_{i \in S_1 \cap S_2} v_i + \sum_{i \in S_2 \setminus S_1} v_i} \right]. \tag{16}
\end{aligned}$$

Here, the first identity holds since the symmetry of \mathbf{p}^* implies that any valuation scenario $\mathbf{v} \in \mathcal{V}$ is a worst-case valuation for \mathbf{p}^* . The second identity assumes that S_1 and S_2 are the assortments corresponding to the first two components of \mathbf{p}' and \mathbf{p}^* , and it follows from the construction of \mathbf{p}' . By choosing $\mathbf{v} \in \mathcal{V}$ such that as many components $i \in S_1 \setminus S_2$ of \mathbf{v} as possible are 0 and as many components $i \in S_2 \setminus S_1$ of \mathbf{v} as possible are 1, we can see that the expression (16) must be strictly negative. Since this implies that the worst-case revenues of \mathbf{p}' are strictly lower than those of \mathbf{p}^* , we obtain the desired contradiction. \square

LEMMA 3. *Under the reference model, any optimal randomization strategy $\mathbf{p} \in \Delta(\mathcal{S})$ satisfies $p_S = 0$ for all $S \in \mathcal{S}$ with $|S| \neq C$.*

Proof of Lemma 3. Note first that any optimal randomization strategy \mathbf{p} has to be feasible and hence place zero probability on any assortment $S \subseteq \mathcal{N}$ violating the cardinality constraint $|S| \leq C$. Assume to the contrary that there is an optimal randomization strategy $\mathbf{p} \in \Delta(\mathcal{S})$ with $p_S > 0$ for some $S \in \mathcal{S}$ satisfying $|S| < C$. We show that \mathbf{p} can be transformed into another optimal randomization strategy $\mathbf{p}' \in \Delta(\mathcal{S})$ satisfying $p'_S = 0$ for all $S \in \mathcal{S}$ with $|S| < C$ and $\mathbf{p}' \neq \mathbf{p}^*$, where \mathbf{p}^* is defined in Lemma 2. Since this contradicts the statement of Lemma 2, we obtain the desired contradiction.

We construct \mathbf{p}' as follows. For any assortment $S \in \mathcal{S}$ with $|S| \leq C$, let $\theta(S) \subseteq \mathcal{N}$ be an assortment satisfying $S \subseteq \theta(S)$ and $|\theta(S)| = C$; for all other assortments S , we define $\theta(S) = S$. We set $p'_T = \sum_{S \in \mathcal{S}: \theta(S)=T} p_S$ for all $T \subseteq \mathcal{N}$ with $|T| = C$, and $p'_T = 0$ otherwise. For later use we note that there are multiple ways to define the mapping θ , and that multiple different probability vectors \mathbf{p}' can be obtained through different θ . We have $\mathbf{p}' \in \Delta(\mathcal{S})$, and the worst-case expected revenues of $\mathbf{p}' \in \Delta(\mathcal{S})$ are at least as large as those of \mathbf{p} since for any valuation scenario $\mathbf{v} \in \mathcal{V}$, we have

$$\begin{aligned}
\sum_{S \in \mathcal{S}} p'_S \cdot R(S, \mathbf{v}) &= \sum_{\substack{S \in \mathcal{S}: \\ |S|=C}} p'_S \cdot R(S, \mathbf{v}) = \sum_{\substack{S \in \mathcal{S}: \\ |S|=C}} p_S \cdot R(S, \mathbf{v}) + \sum_{\substack{S \in \mathcal{S}: \\ |S|<C}} p_S \cdot R(\theta(S), \mathbf{v}) \\
&\geq \sum_{\substack{S \in \mathcal{S}: \\ |S|=C}} p_S \cdot R(S, \mathbf{v}) + \sum_{\substack{S \in \mathcal{S}: \\ |S|<C}} p_S \cdot R(S, \mathbf{v}) = \sum_{S \in \mathcal{S}} p_S \cdot R(S, \mathbf{v}).
\end{aligned}$$

Here, the inequality holds since for every $\mathbf{v} \in \mathcal{V}$, we have

$$R(S, \mathbf{v}) = \frac{\sum_{i \in S} v_i}{v_0 + \sum_{i \in S} v_i} \leq \frac{\sum_{i \in S} v_i + \sum_{i \in \theta(S) \setminus S} v_i}{v_0 + \sum_{i \in S} v_i + \sum_{i \in \theta(S) \setminus S} v_i} = R(\theta(S), \mathbf{v}), \tag{17}$$

where the inequality follows from the fact that $a/b \leq (a+c)/(b+c)$ whenever $0 \leq a < b$ and $c \geq 0$.

If $\mathbf{p}' \neq \mathbf{p}^*$, where \mathbf{p}^* is defined in Lemma 2, then we have the desired contradiction to Lemma 2. Otherwise, if $\mathbf{p}' = \mathbf{p}^*$, then we can exercise our aforementioned flexibility in the choice of θ to derive another optimal randomization strategy $\mathbf{p}'' \neq \mathbf{p}'$ that results in the desired contradiction. \square

Proof of Theorem 2. The first statement follows immediately from Lemmas 1 and 2 and the fact that for any valuation scenario $\mathbf{v} \in \mathcal{V}$, the optimal randomized strategy \mathbf{p}^* places strictly positive probability on assortments S that include products $i \in \mathcal{N}$ with valuation $v_i = 1$. The second statement, on the other hand, follows immediately from Lemmas 2 and 3. \square

Appendix D: Proofs for Section 5

Proof of Observation 1. As per Definition 2, we need to verify that the initial state distribution q is a probability distribution, that for all transition kernels $p \in \mathcal{P}$ and all $x \in \mathcal{X}$ and $a \in \mathcal{A}$, $p(\cdot|x, a)$ is a probability distribution, and that the discount factor γ is contained in the open interval $(0, 1)$.

Since λ is a probability distribution over \mathcal{N}_0 , so is q . The discount factor γ is strictly less than 1 since $\rho_{i0} > 0$ for all $\rho \in \mathcal{U}$ and $i \in \mathcal{N}$. At the same time, γ is strictly positive since $\rho_{i0} < 1$ for at least one $\rho \in \mathcal{U}$ and $i \in \mathcal{N}$. Next, fix any transition kernel $p \in \mathcal{P}$ and state $x \in \mathcal{X}$. For the action $a = \top$, there is precisely one state $x' \in \mathcal{X}$ for which $p(x'|x, a) = 1$, while $p(x''|x, a) = 0$ for all other states $x'' \in \mathcal{X}$. For the action $a = \perp$, we note that for any state $x' \neq 0$, we have

$$0 \leq \rho_{xx'}/\gamma = p(x'|x, a) \leq \sum_{x'' \in \mathcal{N}} p(x''|x, a) = \sum_{x'' \in \mathcal{N}} \rho_{xx''}/\gamma = (1 - \rho_{x0})/\gamma \leq 1,$$

where the first two inequalities hold since $\rho_{xx''} \geq 0$ for all $x'' \neq 0$ and $\gamma > 0$ by construction, the last equality holds since $(\rho_{x0}, \dots, \rho_{xn})$ is a probability distribution over \mathcal{N}_0 , and the last inequality follows from the construction of γ . We thus have $p(x'|x, a) \in [0, 1]$ for all $x' \neq 0$ as well as $\sum_{x'' \in \mathcal{N}} p(x''|x, a) \in [0, 1]$. It then follows directly from our construction in (6) that $p(0|x, a) \in [0, 1]$ as well as $\sum_{x'' \in \mathcal{X}} p(x''|x, a) = 1$. \square

Proof of Theorem 3. For any assortment $S \in \mathcal{S}$, the worst-case expected revenues $R^*(S)$ are given by equation (13) of Désir et al. (2023):

$$\begin{aligned} & \text{minimize} && \sum_{i \in \mathcal{N}_0} \lambda_i \cdot g_i \\ & \text{subject to} && g_i = r_i && \forall i \in S \\ & && g_i = \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j && \forall i \in \mathcal{N} \setminus S \\ & && \rho \in \mathcal{U}, \mathbf{g} \in \mathbb{R}_+^{n+1}. \end{aligned}$$

Similarly, for a policy $\pi_S \in \Pi$, the worst-case expected total discounted reward is given by equation (28) of Nilim and El Ghaoui (2005):

$$\begin{aligned} & \text{minimize} && \sum_{x \in \mathcal{X}} q_x \cdot v_x \\ & \text{subject to} && v_x \geq (1 - \gamma) \cdot r_x + \gamma \cdot v_x && \forall x \in \mathcal{X} : \pi_S(x) = \top \\ & && v_x \geq 0 + \gamma \cdot \sum_{x' \in \mathcal{X}} p(x'|x, \perp) \cdot v_{x'} && \forall x \in \mathcal{X} : \pi_S(x) = \perp \\ & && p \in \mathcal{P}, \mathbf{v} \in \mathbb{R}_+^{n+1}. \end{aligned}$$

Lemma 2 of Nilim and El Ghaoui (2005) implies that there is an optimal solution to this problem that satisfies all inequalities as equalities. We can thus replace all inequalities with equalities, and substituting p with its definition yields

$$\begin{aligned} & \text{minimize} && \sum_{x \in \mathcal{X}} q_x \cdot v_x \\ & \text{subject to} && v_x = r_x && \forall x \in \mathcal{X} : \pi_S(x) = \top \\ & && v_x = \sum_{x' \in \mathcal{X}} \rho_{xx'} \cdot v_{x'} && \forall x \in \mathcal{X} : \pi_S(x) = \perp \\ & && \rho \in \mathcal{U}, \mathbf{v} \in \mathbb{R}_+^{n+1}. \end{aligned}$$

The statement now follows from the fact that $\mathbf{q} = \lambda$ and $\mathcal{X} = \mathcal{N}_0$. \square

Proof of Theorem 4. We first show that a product-wise substitution set in the MC model implies that the ambiguity set of the associated robust MDP in Definition 3 is (x, a) -rectangular. We can then use established results from the robust MDP literature to prove the statement of the theorem.

A robust MDP has an (x, a) -rectangular ambiguity set whenever

$$\mathcal{P} = \{p : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \rightarrow \mathbb{R}_+ : \forall x \in \mathcal{X}, a \in \mathcal{A} \exists p^{xa} \in \mathcal{P} \text{ such that } p(\cdot|x, a) = p^{xa}(\cdot|x, a)\},$$

see Iyengar (2005), Nilim and El Ghaoui (2005) and Wiesemann et al. (2013). In other words, for any selection of conditional distributions p^{xa} , $x \in \mathcal{X}$ and $a \in \mathcal{A}$, we need to show that the composite distribution p satisfying $p(\cdot|x, a) = p^{xa}(\cdot|x, a)$ for all $x \in \mathcal{X}$ and $a \in \mathcal{A}$ is also contained in the ambiguity set \mathcal{P} . Fix any such selection of conditional distributions p^{xa} , $x \in \mathcal{X}$ and $a \in \mathcal{A}$, which by Definition 3 must correspond to a selection of substitution matrices $\rho^{xa} \in \mathcal{U}$, $x \in \mathcal{X}$ and $a \in \mathcal{A}$. Consider the substitution matrix ρ whose i -th row coincides with the i -th row of $\rho^{i\perp}$, $i = 0, \dots, n$. Our definition of product-wise substitution sets guarantees that $\rho \in \mathcal{U}$. Moreover, since $p(\cdot|x, \top)$ does not depend on the substitution matrix, the transition kernel p associated with ρ satisfies $p(\cdot|x, a) = p^{xa}(\cdot|x, a)$ for all $x \in \mathcal{X}$ and $a \in \mathcal{A}$ as desired. We thus conclude that the ambiguity set of the associated robust MDP in Definition 3 is (x, a) -rectangular.

Theorem 4 of Nilim and El Ghaoui (2005) shows that strong duality holds for robust MDPs with (x, a) -rectangular ambiguity sets, that is,

$$\max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right] = \min_{p \in \mathcal{P}} \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right].$$

Theorem 3 of our paper implies that

$$\max_{S \in \mathcal{S}} R^*(S) = \max_{\pi \in \Pi} \min_{p \in \mathcal{P}} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right].$$

Moreover, the same theorem implies that

$$\min_{p \in \mathcal{P}} \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \cdot r(x_t, a_t) \mid x_0 \sim q \right] = \min_{\rho \in \mathcal{U}} \max_{S \in \mathcal{S}} R(S, \rho)$$

since Definition 3 implies the existence of a one-to-one correspondence between transition kernels $p \in \mathcal{P}$ and substitution matrices $\rho \in \mathcal{U}$, and Theorem 3 of our paper can be applied individually to each singleton uncertainty set $\{\rho\}$ constructed from any substitution matrix $\rho \in \mathcal{U}$. We have thus established strong duality for our robust assortment optimization problem, that is,

$$\max_{S \in \mathcal{S}} \min_{\rho \in \mathcal{U}} R(S, \rho) = \min_{\rho \in \mathcal{U}} \max_{S \in \mathcal{S}} R(S, \rho).$$

The right-hand side of this equation bounds from above the worst-case expected revenues achievable by any (deterministic or randomized) assortment strategy. Since this bound is attained by the optimal deterministic assortment on the left-hand side of this equation, we cannot improve the worst-case expected revenues by adopting a randomized assortment strategy. \square

Proof of Proposition 1. For $n = 2$, we can offer three different assortments: $\{1\}$, $\{2\}$ and $\{1, 2\}$. Fix any worst-case substitution matrices $\boldsymbol{\rho}^1 \in \mathcal{U}$ and $\boldsymbol{\rho}^2 \in \mathcal{U}$ for the two assortments $\{1\}$ and $\{2\}$, respectively (they may not be unique). The worst-case expected revenues for the three assortments amount to

$$R(\{1\}, \boldsymbol{\rho}^1) = \lambda_1 \cdot r_1 + \lambda_2 \cdot \rho_{21}^1 \cdot r_1, \quad R(\{2\}, \boldsymbol{\rho}^2) = \lambda_1 \cdot \rho_{12}^2 \cdot r_2 + \lambda_2 \cdot r_2 \quad \text{and} \quad R(\{1, 2\}, \boldsymbol{\rho}) = \lambda_1 \cdot r_1 + \lambda_2 \cdot r_2;$$

note in particular that the expected revenues of the assortment $\{1, 2\}$ do not depend on the realized substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$. We show that the problem is randomization-proof by considering two cases: (i) $R(\{1, 2\}, \boldsymbol{\rho}) \leq \max\{R(\{1\}, \boldsymbol{\rho}^1), R(\{2\}, \boldsymbol{\rho}^2)\}$ and (ii) $R(\{1, 2\}, \boldsymbol{\rho}) > \max\{R(\{1\}, \boldsymbol{\rho}^1), R(\{2\}, \boldsymbol{\rho}^2)\}$.

Case (i). We assume that $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{1\}, \boldsymbol{\rho}^1)$; the case where $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{2\}, \boldsymbol{\rho}^2)$ is symmetric. Since $R(\{1, 2\}, \boldsymbol{\rho}) \leq R(\{1\}, \boldsymbol{\rho}^1)$, we have $\lambda_1 \cdot r_1 + \lambda_2 \cdot r_2 \leq \lambda_1 \cdot r_1 + \lambda_2 \cdot \rho_{21}^1 \cdot r_1$ and thus $\lambda_2 = 0$ or $r_2 \leq \rho_{21}^1 \cdot r_1$. Since only $R(\{2\}, \boldsymbol{\rho})$ depends on $\boldsymbol{\rho}$ if $\lambda_2 = 0$, it is easy to see that the problem is randomization-proof. Assume now that $r_2 \leq \rho_{21}^1 \cdot r_1$. For any substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$, we then have

$$R(\{2\}, \boldsymbol{\rho}) = \lambda_1 \cdot \rho_{12} \cdot r_2 + \lambda_2 \cdot r_2 \leq \lambda_1 \cdot \rho_{12} \cdot \rho_{21}^1 \cdot r_1 + \lambda_2 \cdot \rho_{21}^1 \cdot r_1 \leq \lambda_1 \cdot r_1 + \lambda_2 \cdot \rho_{21}^1 \cdot r_1 = R(\{1\}, \boldsymbol{\rho}^1), \quad (18)$$

where the first inequality holds since $r_2 \leq \rho_{21}^1 \cdot r_1$, while the second inequality is due to the fact that $\rho_{12}, \rho_{21}^1 \leq 1$. Now consider an arbitrary randomized assortment strategy $\alpha \cdot \{1\} + \beta \cdot \{2\} + \gamma \cdot \{1, 2\}$ with $\alpha, \beta, \gamma \geq 0$ and $\alpha + \beta + \gamma = 1$, and let $\boldsymbol{\rho}^r$ be a worst-case substitution matrix for this strategy. We then have that

$$\alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r) + \gamma \cdot R(\{1, 2\}, \boldsymbol{\rho}^r) \leq \alpha \cdot R(\{1\}, \boldsymbol{\rho}^1) + \beta \cdot R(\{1\}, \boldsymbol{\rho}^1) + \gamma \cdot R(\{1\}, \boldsymbol{\rho}^1) = R(\{1\}, \boldsymbol{\rho}^1),$$

where the first inequality holds since $R(\{2\}, \boldsymbol{\rho}^r) \leq R(\{1\}, \boldsymbol{\rho}^1)$ by (18) and $R(\{1, 2\}, \boldsymbol{\rho}^r) \leq R(\{1\}, \boldsymbol{\rho}^1)$ by assumption. We thus conclude that the problem is randomization-proof.

Case (ii). We prove this case by contradiction. We first show that if a problem instance is randomization-receptive, then it is optimal to randomize between the singleton assortments $\{1\}$ and $\{2\}$. We then derive a necessary and sufficient condition for the randomization between $\{1\}$ and $\{2\}$ to be beneficial. Finally, we show that this condition contradicts our assumption that $R(\{1, 2\}, \boldsymbol{\rho}) > \max\{R(\{1\}, \boldsymbol{\rho}^1), R(\{2\}, \boldsymbol{\rho}^2)\}$.

Assume that the problem is randomization-receptive, and let $\alpha \cdot \{1\} + \beta \cdot \{2\} + \gamma \cdot \{1, 2\}$ with $\alpha, \beta \geq 0, \gamma > 0$ and $\alpha + \beta + \gamma = 1$ be an optimal randomization strategy. Denote by $\boldsymbol{\rho}^r$ one of its worst-case substitution matrices. We then have $\alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r) + \gamma \cdot R(\{1, 2\}, \boldsymbol{\rho}^r) > R(\{1, 2\}, \boldsymbol{\rho}^r)$, and hence $(1 - \gamma) \cdot R(\{1, 2\}, \boldsymbol{\rho}^r) < \alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r)$. This implies that

$$\begin{aligned} \alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r) + \gamma \cdot R(\{1, 2\}, \boldsymbol{\rho}^r) &< \left[1 + \frac{\gamma}{1 - \gamma}\right] \cdot (\alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r)) \\ &= \frac{\alpha}{\alpha + \beta} \cdot R(\{1\}, \boldsymbol{\rho}^r) + \frac{\beta}{\alpha + \beta} \cdot R(\{2\}, \boldsymbol{\rho}^r), \end{aligned}$$

where the inequality uses the fact that $(1 - \gamma) \cdot R(\{1, 2\}, \boldsymbol{\rho}^r) < \alpha \cdot R(\{1\}, \boldsymbol{\rho}^r) + \beta \cdot R(\{2\}, \boldsymbol{\rho}^r)$, while the equality holds since $\gamma = 1 - \alpha - \beta$. Since $R(\{1, 2\}, \boldsymbol{\rho})$ does not depend on $\boldsymbol{\rho}$, we have $\boldsymbol{\rho}^r \in \arg \min_{\boldsymbol{\rho} \in \mathcal{U}} \frac{\alpha}{\alpha + \beta} \cdot R(\{1\}, \boldsymbol{\rho}) + \frac{\beta}{\alpha + \beta} \cdot R(\{2\}, \boldsymbol{\rho})$. It is therefore optimal to set $\gamma = 0$ and randomize only between the assortments $\{1\}$ and $\{2\}$.

The problem of randomizing between the assortments $\{1\}$ and $\{2\}$ can be written as

$$\begin{aligned} & \text{maximize}_{\beta \in [0,1]} \left[\min_{\rho \in \mathcal{U}} \{ \beta \cdot R(\{1\}, \rho) + (1 - \beta) \cdot R(\{2\}, \rho) \} \right] \\ &= \text{maximize}_{\beta \in [0,1]} \left[\min_{\rho \in \mathcal{U}} \{ \lambda_1 \cdot r_1 + \lambda_2 \cdot r_2 + \beta \cdot \lambda_2 (\rho_{21} \cdot r_1 - r_2) + (1 - \beta) \cdot \lambda_1 \cdot (\rho_{12} \cdot r_2 - r_1) \} \right]. \end{aligned}$$

For the problem to be randomization-receptive, it thus needs to satisfy

$$\begin{aligned} & \text{maximize}_{\beta \in [0,1]} \left[\min_{\rho \in \mathcal{U}} \{ \lambda_1 \cdot r_1 + \lambda_2 \cdot r_2 + \beta \cdot \lambda_2 \cdot (\rho_{21} \cdot r_1 - r_2) + (1 - \beta) \cdot \lambda_1 \cdot (\rho_{12} \cdot r_2 - r_1) \} \right] \\ & \quad > \lambda_1 \cdot \rho_{12} \cdot r_2 + \lambda_2 \cdot r_2 \\ \iff & \exists \beta \in (0, 1) \text{ such that } \beta \cdot \lambda_2 \cdot (\rho_{21} \cdot r_1 - r_2) + (1 - \beta) \cdot \lambda_1 \cdot (\rho_{12} \cdot r_2 - r_1) > 0 \quad \forall \rho \in \mathcal{U}, \end{aligned} \quad (19)$$

where we use our case assumption that the assortment $\{1, 2\}$ generates the highest worst-case expected revenues. We now show that (19) contradicts the assumption that $R(\{1, 2\}, \rho) > \max\{R(\{1\}, \rho^1), R(\{2\}, \rho^2)\}$.

By our case assumption, we have $\lambda_1 \cdot r_1 + \lambda_2 \cdot r_2 > \lambda_1 \cdot r_1 + \lambda_2 \cdot \rho_{21}^1 \cdot r_1$ and $\lambda_1 \cdot r_1 + \lambda_2 \cdot r_2 > \lambda_2 \cdot r_2 + \lambda_1 \cdot \rho_{12}^2 \cdot r_2$. These inequalities imply that $\lambda_1, \lambda_2 > 0$ as well as $r_2 > \rho_{21}^1 \cdot r_1$ and $r_1 > \rho_{12}^2 \cdot r_2$. For notational convenience, we introduce the constants a, b and c defined via:

$$r_2 = a \cdot r_1 \text{ for some } a \in (\rho_{21}^1, 1/\rho_{12}^2), \quad \lambda_2 = b \cdot \lambda_1 \text{ for some } b \in \mathbb{R}_{++}, \quad 1 - \beta = c \cdot \beta \text{ for some } c \in \mathbb{R}_{++}.$$

Using this notation, we can re-express the inequality in (19) as

$$\begin{aligned} & \beta \cdot \lambda_2 \cdot (\rho_{21} \cdot r_1 - r_2) + (1 - \beta) \cdot \lambda_1 \cdot (\rho_{12} \cdot r_2 - r_1) > 0 \\ \iff & \beta \cdot b \cdot \lambda_1 \cdot (\rho_{21} \cdot r_1 - a \cdot r_1) + c \cdot \beta \cdot \lambda_1 \cdot (\rho_{12} \cdot a \cdot r_1 - r_1) > 0 \\ \iff & b \cdot \rho_{21} - b \cdot a > c - c \cdot a \cdot \rho_{12}. \end{aligned}$$

Here, the first equivalence results from the definitions of a, b and c , whereas the second equivalence follows from dividing the expressions by $\beta \lambda_1 r_1 > 0$. Since this inequality has to hold for all $\rho \in \mathcal{U}$, it has to hold in particular for ρ^1 and ρ^2 , which respectively result in the conditions

$$\begin{aligned} b \cdot \rho_{21}^1 - b \cdot a > c - c \cdot a \cdot \rho_{12}^1 & \implies b \cdot \rho_{21}^1 - b \cdot a > c - c \cdot a \\ & \implies b \cdot (\rho_{21}^1 - a) > c \cdot (1 - a) & \implies a > 1 \\ \text{and } b \cdot \rho_{21}^2 - b \cdot a > c - c \cdot a \cdot \rho_{12}^2 & \implies b - b \cdot a > c - c \cdot a \cdot \rho_{12}^2 \\ & \implies b \cdot (1 - a) > c \cdot (1 - a \cdot \rho_{12}^2) & \implies a < 1. \end{aligned}$$

Here, the respective first implications hold since $\rho_{12}^1, \rho_{21}^2 \in [0, 1]$, and the last ones follow from the fact that $a \in (\rho_{21}^1, 1/\rho_{12}^2)$ and $b, c > 0$. This is a contradiction, and we thus conclude that the problem must be randomization-proof under case (ii) as well. \square

Proof of Observation 2. Consider an instance with n products where the product prices are $\mathbf{r} = (0, 4.66, 1, 10, 1, \dots, 1)$, the initial choice probabilities are $\boldsymbol{\lambda} = (0, 0.37, 0.62, 0.01, 0, \dots, 0)$, and the uncertainty set is

$$\mathcal{U} = \left\{ \left(\begin{array}{cccccccc} 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0.01 & 0 & 0.05 & 0.94 & 0 & \dots & 0 \\ 0.26 & 0.69 & 0 & 0.05 & 0 & \dots & 0 \\ 0.90 & 0.05 & 0.05 & 0 & 0 & \dots & 0 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \end{array} \right), \left(\begin{array}{cccccccc} 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0.90 & 0 & 0.05 & 0.05 & 0 & \dots & 0 \\ 0.01 & 0.44 & 0 & 0.55 & 0 & \dots & 0 \\ 0.01 & 0.94 & 0.05 & 0 & 0 & \dots & 0 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \end{array} \right) \right\}.$$

First, one readily verifies that \mathcal{U} does not have product-wise substitution sets. Since the products $4, \dots, n$

Table 10 Expected revenues of different assortments (worst-case scenarios highlighted in bold).

Assortment	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	{1,2,3}	$0.374 \cdot \{3\} + 0.626 \cdot \{1,3\}$
Scenario 1	3.72	0.66	7.93	2.35	4.13	4.22	2.44	5.68
Scenario 2	4.62	0.64	4.02	2.39	6.57	0.92	2.44	5.61

are never purchased and hence do not affect the revenues, it is optimal to offer assortments containing only (subsets of) the first three products. Table 10 lists the expected revenues of all deterministic assortments as well as one randomized assortment under both scenarios in \mathcal{U} . The table shows that the randomized assortment outperforms the deterministic assortments in terms of worst-case expected revenues. We thus conclude that the problem instance is randomization-receptive. \square

Proof of Observation 3. We show that $R_{\text{det}}^*(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \cdot \min\{r_i : i \in \mathcal{N}\}$ and $R_{\text{rand}}^*(\mathcal{U}) \leq \sum_{i \in \mathcal{N}} \lambda_i \cdot \max\{r_i : i \in \mathcal{N}\}$, which together imply the statement. To see that $R_{\text{det}}^*(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \cdot \min\{r_i : i \in \mathcal{N}\}$, consider the deterministic assortment $S = \mathcal{N}$. The worst-case expected revenues of this assortment amount to $\sum_{i \in \mathcal{N}} \lambda_i \cdot r_i \geq \sum_{i \in \mathcal{N}} \lambda_i \cdot \min\{r_i : i \in \mathcal{N}\}$, and hence $R_{\text{det}}^*(\mathcal{U}) \geq \sum_{i \in \mathcal{N}} \lambda_i \cdot \min\{r_i : i \in \mathcal{N}\}$. To see that $R_{\text{rand}}^*(\mathcal{U}) \leq \sum_{i \in \mathcal{N}} \lambda_i \cdot \max\{r_i : i \in \mathcal{N}\}$, on the other hand, note that any customer will exercise the no-purchase option with a probability of at least $\lambda_0 = 1 - \sum_{i \in \mathcal{N}} \lambda_i$. Thus, even if a randomized assortment strategy would achieve expected revenues of $\max\{r_i : i \in \mathcal{N}\}$ whenever a customer purchases any product, its worst-case expected revenues are bounded above by $\sum_{i \in \mathcal{N}} \lambda_i \cdot \max\{r_i : i \in \mathcal{N}\}$. \square

Proof of Theorem 5. In view of the first statement, fix any $n \geq 3$ and $C \leq n - 2$. We construct the following instance of the cardinality-constrained robust MC problem. The product prices are $r_1 = 0$ and $r_i = 1$, $i \in \{2, \dots, n\}$, the initial choice probabilities satisfy $\lambda_1 = 1$ and $\lambda_i = 0$, $i \neq 1$, and the uncertainty set \mathcal{U} contains all substitution matrices $\boldsymbol{\rho}$ satisfying $\boldsymbol{\rho}_0 = (1, 0, \dots, 0)$, $\boldsymbol{\rho}_1 = (0, \boldsymbol{\rho}')$ where $\boldsymbol{\rho}'$ is any element of the probability simplex in \mathbb{R}^n , as well as $\boldsymbol{\rho}_i$ being any element of the probability simplex in \mathbb{R}^{n+1} , $i \in \{2, \dots, n\}$. Note that by construction, this uncertainty set exhibits product-wise substitution sets.

The worst-case expected revenues of every deterministic assortment $S \in \mathcal{S}$ satisfying $|S| \leq C$ are 0. This is clearly the case if $1 \in S$. If $1 \notin S$, fix any product $i \in \mathcal{N} \setminus (S \cup \{1\})$, which exists since $C \leq n - 2$. For any substitution matrix $\boldsymbol{\rho}^i$ satisfying $\rho_{1i}^i = \rho_{i0}^i = 1$, we have $R(S, \boldsymbol{\rho}^i) = 0$.

We now consider any randomized assortment strategy that places a strictly positive probability on every assortment $S \in \mathcal{S} \setminus \{1\}$ satisfying $|S| \leq C$ and that places zero probability on all other assortments. Fix any substitution matrix $\boldsymbol{\rho} \in \mathcal{U}$. Since $\rho_{1i} \geq 1/n$ for at least one $i \in \{2, \dots, n\}$ and since this product i will be offered with strictly positive probability, the randomized assortment strategy must generate strictly positive expected revenues under the substitution matrix $\boldsymbol{\rho}$. Since $\boldsymbol{\rho}$ was arbitrary, the worst-case expected revenues of the randomized assortment strategy must be strictly positive as well.

As for the second statement, consider an instance of the randomized robust assortment optimization problem under an MC model with product-wise substitution sets, $n \geq 2$ products with $r_i = 1$ and $\lambda_i = 1/n$ for all $i \in \mathcal{N}$, as well as a cardinality constraint $|S| \leq C$ with $C \in \{1, \dots, n - 2\}$. We consider the product-wise substitution sets

$$\mathcal{U}_i = \left\{ \frac{1}{n-C} \cdot (\mathbf{e}_0 + \sum_{k \in \mathcal{K}} \mathbf{e}_k) : \mathcal{K} \subseteq \mathcal{N} \setminus \{i\}, |\mathcal{K}| = n - C - 1 \right\} \quad \forall i \in \mathcal{N},$$

where by a slight abuse of notation the indices of $\mathbf{e}_i \in \mathbb{R}^{n+1}$ start from 0 so as to match the product labels. The uncertainty set \mathcal{U}_0 associated with product 0 is the singleton set $\{(1, 0, \dots, 0)\}$.

Under this setting, we can follow the same steps as in the proof of [Lemma 2](#) to show that the randomized assortment strategy $\mathbf{p}^* \in \Delta(\mathcal{S})$, $p_S^* = 1/|\{S \in \mathcal{S} : |S| = C\}|$ if $|S| = C$ and $p_S^* = 0$ otherwise, strictly outperforms any other feasible randomized assortment strategy \mathbf{p} that only places positive probability $p_S > 0$ on assortments S with $|S| = C$. The only aspect that needs to be verified is equation (16), which holds since

$$\begin{aligned} & \min_{\boldsymbol{\rho} \in \mathcal{U}} \epsilon \cdot (R(S_1, \boldsymbol{\rho}) - R(S_2, \boldsymbol{\rho})) \\ = & \min_{\boldsymbol{\rho} \in \mathcal{U}} \epsilon \cdot \left[\sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^{S_1} - \sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^{S_2} \right] \quad \text{with } g_i^{S_1} = \begin{cases} r_i & \text{if } i \in S_1 \\ \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j^{S_1} & \text{if } i \in \mathcal{N} \setminus S_1 \end{cases} \\ & \quad \text{and } g_i^{S_2} = \begin{cases} r_i & \text{if } i \in S_2 \\ \sum_{j \in \mathcal{N}} \rho_{ij} \cdot g_j^{S_2} & \text{if } i \in \mathcal{N} \setminus S_2 \end{cases} \\ = & \min_{\boldsymbol{\rho} \in \mathcal{U}} \epsilon \cdot \left[\sum_{i \in \mathcal{N} \setminus S_1} \lambda_i \cdot g_i^{S_1} - \sum_{i \in \mathcal{N} \setminus S_2} \lambda_i \cdot g_i^{S_2} \right] < 0. \end{aligned}$$

Here, the first identity is due to the definition of $R(S_1, \boldsymbol{\rho})$ and $R(S_2, \boldsymbol{\rho})$, while the second equality holds since both S_1 and S_2 contain C products and the prices r_i and initial probabilities λ_i of all products are the same. The inequality, finally, is due to the fact that we can choose $\boldsymbol{\rho}_i \in \mathcal{U}_i$, $i \in \mathcal{N}$, whose associated choice of \mathcal{K} satisfies $\mathcal{K} \cap S_2 \neq \emptyset$ whenever $i \in S_1 \setminus S_2$ and $\mathcal{K} \cap S_1 = \emptyset$ otherwise. This choice of $\boldsymbol{\rho}$ ensures that (i) when S_2 is offered, customers demanding products in $\mathcal{N} \setminus S_2$ subsequently consider products in S_2 and hence generate revenues, whereas (ii) when S_1 is offered, customers demanding products in $\mathcal{N} \setminus S_1$ will not consider any products in S_1 , and hence no extra revenues accrue. This explains why the difference is strictly negative.

Having established the analogue of [Lemma 2](#) for our current setting, we can now apply the steps in the proof of [Lemma 3](#) to show that any optimal randomization strategy $\mathbf{p} \in \Delta(\mathcal{S})$ satisfies $p_S = 0$ for all $S \in \mathcal{S}$ with $|S| \neq C$. To this end, the only aspect that needs to be verified is

$$R(S, \boldsymbol{\rho}) = \sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^S \leq \sum_{i \in \mathcal{N}} \lambda_i \cdot g_i^{\theta(S)} = R(\theta(S), \boldsymbol{\rho}),$$

where $\theta(S)$ is any assortment of cardinality C that contains S . The inequality holds since for all $i \in \theta(S) \setminus S$, we have $g_i^{\theta(S)} = 1 \geq g_i^S$, which in turn implies that $g_i^{\theta(S)} \geq g_i^S$ for all $i \in \mathcal{N} \setminus \theta(S)$. \square

Table 11 Expected revenues of different assortments consisting of only products 1 and 2.

Scenario	{1}	{2}	{1, 2}	$\frac{1}{3} \cdot \{2\} + \frac{2}{3} \cdot \{1, 2\}$
1	1	2	1	1.33
2	1	1	1.5	1.33

Appendix E: Proofs for Section 6

Proof of Observation 4. Consider an assortment optimization problem under the PR model as below defined next. The product prices are $r_1 = 1$, $r_2 = 2$ and $r_i = 1$ for $i \geq 3$. The preference rankings are

$$1 \rightarrow 2 \rightarrow 0 \rightarrow 3 \rightarrow 4 \rightarrow \dots \rightarrow n, \quad 2 \rightarrow 1 \rightarrow 0 \rightarrow 3 \rightarrow 4 \rightarrow \dots \rightarrow n, \quad 1 \rightarrow 0 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow \dots \rightarrow n.$$

We assume that the ambiguity set Λ contains the two occurrence probability scenarios $\lambda^1 = (1, 0, 0)$ and $\lambda^2 = (0, 0.5, 0.5)$. Note that the products 3 to n are consistently ranked lower than product 0 in all preference rankings. Therefore, we can restrict our attention to assortments containing the products 1 and 2 only. The expected revenues of the corresponding deterministic assortments as well as one randomized assortment are listed in Table 11. As the table shows, the randomized assortment generates strictly higher worst-case expected revenues than the best deterministic assortment, which concludes the proof. \square

Proof of Observation 5. The bound for a generic number of products n can be derived using similar arguments as in the proof of Observation 3; we therefore only prove the bound for $n = 2$ products. To avoid trivial cases, we assume that $r_1 \neq r_2$ and $r_1, r_2 > 0$. Without loss of generality, we further assume that $r_1 > r_2$. In this case, including product 1 into an assortment weakly increases the expected revenues under any occurrence probability scenario, and therefore we can restrict ourselves to deterministic and randomized assortments that include product 1. We then consider the preference rankings

$$\begin{aligned} 1: & \quad 1 \rightarrow 0 & \quad \text{with probability } \lambda_1 \\ 2: & \quad 2 \rightarrow 1 \rightarrow 0 & \quad \text{with probability } \lambda_2 \\ 3: & \quad 2 \rightarrow 0 & \quad \text{with probability } \lambda_3, \end{aligned}$$

where $\lambda_1, \lambda_2, \lambda_3 \geq 0$ and $\lambda_1 + \lambda_2 + \lambda_3 = 1$. Note that we do not consider any preference rankings that contain the no-purchase option 0 as highest ranked choice; this is without loss of generality as such a preference ranking affects every (deterministic and randomized) assortment in the same way. We now compute the expected revenues of the deterministic and randomized assortments of interest:

$$\begin{aligned} \{1\}: & \quad (\lambda_1 + \lambda_2) \cdot r_1 = (1 - \lambda_3) \cdot r_1 \\ \{1, 2\}: & \quad \lambda_1 \cdot r_1 + (\lambda_2 + \lambda_3) \cdot r_2 = r_2 + \lambda_1 \cdot (r_1 - r_2) \\ \alpha \cdot \{1\} + \beta \cdot \{1, 2\}: & \quad \alpha \cdot (1 - \lambda_3) \cdot r_1 + \beta \cdot [r_2 + \lambda_1 \cdot (r_1 - r_2)] \end{aligned}$$

Fix an optimal randomization strategy (α^*, β^*) as well as the worst-case distributions (not necessarily unique) $\lambda^1 = (\lambda_1^1, \lambda_2^1, \lambda_3^1) \in \mathcal{U}$, $\lambda^{12} = (\lambda_1^{12}, \lambda_2^{12}, \lambda_3^{12}) \in \mathcal{U}$ and $\lambda^r = (\lambda_1^r, \lambda_2^r, \lambda_3^r) \in \mathcal{U}$ for the deterministic assortments $\{1\}$ and $\{1, 2\}$ as well as the randomized assortment $\alpha^* \cdot \{1\} + \beta^* \cdot \{1, 2\}$, respectively.

The benefit of randomization can now be computed as

$$R = \frac{\alpha^* \cdot (1 - \lambda_3^r) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^r \cdot (r_1 - r_2)]}{\max \{(1 - \lambda_3^1) \cdot r_1, r_2 + \lambda_1^{12} \cdot (r_1 - r_2)\}}.$$

Assume first that $(1 - \lambda_3^1) \cdot r_1 \geq r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$ in the denominator of R . We then have

$$\begin{aligned} R &= \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} \leq \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} \\ &= 1 + \frac{-\beta^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} = 1 + \beta^* \cdot \frac{-(1 - \lambda_3^1) \cdot r_1 + r_2 + \lambda_1^1 \cdot (r_1 - r_2)}{(1 - \lambda_3^1) \cdot r_1} \\ &\leq 1 + \beta^* \cdot \frac{-(1 - \lambda_3^1) \cdot r_1 + r_2 + (1 - \lambda_3^1) \cdot (r_1 - r_2)}{(1 - \lambda_3^1) \cdot r_1} = 1 + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{(1 - \lambda_3^1) \cdot r_1} \leq 1 + \beta^* \leq 2, \end{aligned}$$

where the first equality is due to $(1 - \lambda_3^1) \cdot r_1 \geq r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$, the first inequality holds since the occurrence probabilities λ^r minimize the expected revenues of the randomized assortment, the second equality follows from $\alpha^* + \beta^* = 1$, the second inequality holds since $\lambda_1^1, \lambda_2^1, \lambda_3^1 \geq 0$ and $\lambda_1^1 + \lambda_2^1 + \lambda_3^1 = 1$, and the third inequality follows from the fact that $(1 - \lambda_3^1) \cdot r_1 \geq r_2 + \lambda_1^{12} \cdot (r_1 - r_2) \geq r_2$ since $r_1 > r_2$ by assumption.

Assume next that $(1 - \lambda_3^1) \cdot r_1 < r_2 + \lambda_1^{12} \cdot (r_1 - r_2)$ in the denominator of R . We then have

$$\begin{aligned} R &= \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \leq \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \\ &= \frac{\alpha^* \cdot (1 - \lambda_3^1) \cdot r_1 + \beta^* \cdot [r_2 + \lambda_1^1 \cdot (r_1 - r_2)]}{(1 - \lambda_3^1) \cdot r_1} \cdot \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \\ &\leq \left(1 + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{(1 - \lambda_3^1) \cdot r_1} \right) \cdot \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} = \frac{(1 - \lambda_3^1) \cdot r_1}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} + \beta^* \cdot \frac{\lambda_3^1 \cdot r_2}{r_2 + \lambda_1^{12} \cdot (r_1 - r_2)} \leq 2, \end{aligned}$$

where the calculations follow the same ideas as in the previous case. Note that our derivations assume that $(1 - \lambda_3^1) \cdot r_1 \neq 0$. Indeed, if $(1 - \lambda_3^1) \cdot r_1 = 0$, then $\lambda_3^1 = 1$ since $r_1 > 0$ by assumption. In that case, however, the problem is readily verified to be randomization-proof. \square

We now show that the two product bound of [Observation 5](#) is indeed attained asymptotically.

EXAMPLE 2. Consider the unconstrained robust PR problem with two products, $r_1 = M$ (with M being a large positive number) and $r_2 = 1$ as well as the preference rankings

1:	1 \rightarrow 0	with probability λ_1
2:	2 \rightarrow 1 \rightarrow 0	with probability λ_2
3:	2 \rightarrow 0	with probability λ_3 .

We assume that the ambiguity set \mathcal{U} contains the two occurrence probability scenarios $\lambda^1 = (0, 1, 0)$ and $\lambda^2 = (1/M, 0, 1 - 1/M)$. Following the same reasoning as in the proof of [Observation 5](#), we can without loss of generality restrict ourselves to the consideration of the following assortments with the associated worst-case expected revenues:

$$\begin{aligned} \{1\}: & \quad \min \{M, 1\} = 1 \\ \{1, 2\}: & \quad \min \{1, 2 - 1/M\} = 1 \\ \alpha \cdot \{1\} + \beta \cdot \{1, 2\}: & \quad \min \{\alpha \cdot M + \beta, 1 + \beta \cdot (1 - 1/M)\} \end{aligned}$$

Clearly, $R_{\det}^*(\mathcal{V}) = 1$ in this instance. The worst-case expected revenues of the randomized assortment strategy $(\alpha^*, \beta^*) = \left(\frac{1-1/M}{M-1/M}, \frac{M-1}{M-1/M} \right)$, on the other hand, evaluate to

$$\min \left\{ \frac{1-1/M}{M-1/M} \cdot M + \frac{M-1}{M-1/M}, 1 + \frac{M-1}{M-1/M} \cdot (1-1/M) \right\} = \frac{2M-2}{M-1/M},$$

and the last expression evidently converges to 2 as $M \rightarrow \infty$.

Proof of Theorem 6. Consider the cardinality-constrained PR problem with product prices $r_1 = \dots = r_n = 1$ and the n preference rankings

$$1 \rightarrow 0 \rightarrow \dots, \quad 2 \rightarrow 0 \rightarrow \dots, \quad \dots, \quad n \rightarrow 0 \rightarrow \dots,$$

where the preferences after product 0 are omitted since they do not affect the revenues. We fix the ambiguity set \mathcal{U} that contains all probability distributions $\boldsymbol{\lambda} \in \mathbb{R}^n$ contained in the n -dimensional probability simplex. Any deterministic assortment $S \in \mathcal{S}$ satisfies $|S| \leq C$ and thus generates zero revenues under the Dirac distribution that places all probability mass on a preference ranking containing a product $i \in \mathcal{N} \setminus S$ as highest ranked product. Consider now any randomized assortment strategy that places a strictly positive probability on every assortment $S \in \mathcal{S}$. Since every product is offered with a strictly positive probability, this strategy generates strictly positive expected revenues under any occurrence probability scenario $\boldsymbol{\lambda} \in \mathcal{U}$. We thus conclude that the worst-case expected revenues of this strategy are also strictly positive. \square

Proof of Theorem 7. In view of the first statement, fix any optimal randomization strategy \boldsymbol{p}^* for the robust PR problem. If at most $K+1$ components of \boldsymbol{p}^* are positive, then there is nothing to prove. Otherwise, note that the worst-case revenues of \boldsymbol{p}^* evaluate to

$$\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} p_S^* \cdot R(S, \boldsymbol{\lambda}) = \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} \sum_{k \in \mathcal{K}} p_S^* \cdot \lambda_k \cdot R_k(S) = \min_{\boldsymbol{\lambda} \in \Lambda} \boldsymbol{\lambda}^\top \left[\sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{R}(S) \right], \quad (20)$$

where $\boldsymbol{R}(S) = (R_1(S), \dots, R_K(S))^\top \in \mathbb{R}^K$. Here, the summation $\sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{R}(S)$ can be interpreted as a convex combination of the (potentially exponentially many) expected revenues vectors $\boldsymbol{R}(S) \in \mathbb{R}^K$, $S \in \mathcal{S}$. Carathéodory's theorem then implies that there are alternative weights q_S^* , $S \in \mathcal{S}$, such that at most $K+1$ components of \boldsymbol{q}^* are positive and that satisfy

$$\sum_{S \in \mathcal{S}} q_S^* \cdot \boldsymbol{R}(S) = \sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{R}(S).$$

Substituting the expression on the left-hand side of this identity into equation (20), we conclude that

$$\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} q_S^* \cdot R(S, \boldsymbol{\lambda}) = \min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} p_S^* \cdot R(S, \boldsymbol{\lambda}),$$

which implies that the parsimonious randomization strategy \boldsymbol{q}^* attains the same objective value as the optimal randomization strategy \boldsymbol{p}^* .

In view of the second statement, assume that the uncertainty set is of the form

$$\Lambda = \boldsymbol{F}\Xi \quad \text{with} \quad \Xi = \{\boldsymbol{\xi} \in \mathbb{R}_+^m : \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b}\} \subseteq \mathbb{R}_+^m.$$

Fix any optimal randomization strategy \boldsymbol{p}^* for the randomized robust PR problem. If at most $m+1$ components of \boldsymbol{p}^* are positive, then there is nothing to prove. Otherwise, similar to (20), the worst-case revenues of \boldsymbol{p}^* evaluate to

$$\min_{\boldsymbol{\lambda} \in \Lambda} \sum_{S \in \mathcal{S}} p_S^* \cdot R(S, \boldsymbol{\lambda}) = \min_{\boldsymbol{\xi} \in \Xi} (\boldsymbol{F}\boldsymbol{\xi})^\top \left[\sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{R}(S) \right] = \min_{\boldsymbol{\xi} \in \Xi} \boldsymbol{\xi}^\top \left[\sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{F}^\top \boldsymbol{R}(S) \right].$$

Here, the summation $\sum_{S \in \mathcal{S}} p_S^* \cdot \boldsymbol{F}^\top \boldsymbol{R}(S)$ can be interpreted as a convex combination of the (potentially exponentially many) vectors $\boldsymbol{F}^\top \boldsymbol{R}(S) \in \mathbb{R}^m$, $S \in \mathcal{S}$. Applying Carathéodory's theorem again shows that there are alternative weights q_S^* , $S \in \mathcal{S}$, with at most $m+1$ components positive such that \boldsymbol{q}^* attains the same objective value as the optimal randomization strategy \boldsymbol{p}^* . This concludes the proof. \square

Appendix F: Supplementary Numerical Results for Section 7

Tables 12 and 13 provide supplementary information on the runtimes as well as the performance of the different problem formulations for the MC choice model.

n	K	deterministic		randomized	
		exact runtime	exact objective	runtime	heuristic objective runtime
20	5	0.22s	0.91%	1.02s	0.90% 0.50s
20	8	0.73s	1.29%	7.32s	1.28% 5.96s
20	10	2.22s	1.35%	31.32s	1.33% 29.29s
20	12	8.17s	1.48%	165.12s	1.47% 150.10s
30	5	0.60s	0.65%	2.14s	0.65% 1.35s
30	8	1.65s	0.81%	12.18s	0.81% 9.78s
30	10	4.29s	0.96%	55.10s	0.95% 47.74s
30	12	14.82s	1.05%	247.61s	1.05% 230.06s
40	5	1.07s	0.45%	3.37s	0.44% 3.21s
40	8	2.79s	0.60%	18.19s	0.59% 16.09s
40	10	6.09s	0.63%	75.93s	0.62% 71.59s
40	12	20.08s	0.73%	331.86s	0.72% 303.40s
50	5	1.77s	0.37%	4.68s	0.36% 6.54s
50	8	5.65s	0.44%	29.14s	0.44% 30.08s
50	10	11.10s	0.47%	126.39s	0.47% 118.87s
50	12	29.71s	0.55%	543.70s	0.55% 508.67s
60	5	3.11s	0.29%	7.21s	0.28% 13.84s
60	8	7.86s	0.32%	39.07s	0.32% 47.93s
60	10	19.32s	0.37%	180.43s	0.37% 176.34s
60	12	41.94s	0.41%	657.93s	0.41% 642.21s
70	5	3.68s	0.22%	8.55s	0.21% 25.21s
70	8	10.26s	0.26%	48.36s	0.26% 70.98s
70	10	25.73s	0.32%	203.54s	0.32% 227.00s
70	12	54.15s	0.32%	792.60s	0.32% 755.90s
80	5	6.56s	0.20%	11.66s	0.19% 43.60s
80	8	15.46s	0.23%	58.25s	0.23% 96.88s
80	10	35.85s	0.27%	254.84s	0.26% 300.97s
80	12	75.24s	0.27%	1,118.74s	0.27% 1,165.78s
90	5	8.66s	0.17%	15.07s	0.16% 83.25s
90	8	27.11s	0.17%	81.19s	0.17% 158.98s
90	10	57.35s	0.22%	366.83s	0.22% 451.94s
90	12	89.78s	0.22%	1,109.80s	0.22% 1,105.29s
100	5	17.66s	0.17%	23.08s	0.17% 142.46s
100	8	36.31s	0.17%	111.31s	0.17% 260.03s
100	10	68.60s	0.18%	404.69s	0.17% 538.53s
100	12	107.03s	0.21%	1,244.29s	0.21% 1,412.64s

Table 12 Deterministic vs. randomized solution of the unconstrained MC problem under general substitution sets. All percentages are reported relative to the worst-case expected revenues of the optimal deterministic robust assortment.

n	K	deterministic		randomized		
		exact runtime	objective	exact runtime	objective	heuristic runtime
20	5	1.26s	4.51%	6.23s	4.45%	1.30s
20	8	5.08s	7.53%	35.01s	7.48%	18.47s
20	10	14.11s	8.70%	144.44s	8.65%	100.67s
20	12	43.62s	9.20%	556.59s	9.16%	474.12s
30	5	6.04s	6.35%	35.69s	6.29%	4.06s
30	8	35.29s	9.93%	200.57s	9.86%	49.72s
30	10	98.11s	12.01%	565.06s	11.92%	299.34s
30	12	233.42s	13.48%	2,112.93s	13.41%	1,531.79s
40	5	18.21s	7.85%	131.25s	7.76%	11.94s
40	8	174.41s	12.52%	1,042.42s	12.41%	137.71s
40	10	498.01s	15.20%	2,898.64s	15.10%	684.67s
40	12	1,404.03s	16.94%	8,288.72s	16.83%	3,852.12s
50	5	44.81s	9.17%	458.82s	9.06%	33.16s
50	8	572.67s	14.82%	3,955.41s	14.70%	282.57s
50	10	2,260.97s	17.45%	9,981.55s	17.33%	1,481.38s
50	12	3,608.87s	17.80%	17,054.40s	17.70%	6,706.03s
60	5	80.90s	9.99%	1,073.07s	9.90%	61.95s
60	8	1,441.70s	15.90%	9,103.87s	15.78%	519.43s
60	10	4,789.94s	18.22%	21,139.40s	18.11%	2,592.26s
60	12	6,440.95s	18.27%	27,527.50s	18.17%	10,547.90s
70	5	141.82s	11.40%	2,093.76s	11.30%	132.29s
70	8	2,784.57s	16.60%	19,390.60s	16.47%	928.64s
70	10	6,113.36s	17.82%	32,640.60s	17.68%	3,538.27s
70	12	6,237.76s	16.93%	33,286.90s	16.85%	13,432.00s
80	5	220.94s	11.92%	3,198.93s	11.80%	222.90s
80	8	4,132.22s	17.56%	30,642.75s	17.43%	1,431.79s
80	10	6,352.38s	16.79%	34,369.30s	16.69%	4,526.13s
80	12	10,617.26s	17.70%	28,164.65s	17.65%	15,879.15s
90	5	367.60s	12.94%	5,059.18s	12.83%	399.70s
90	8	4,020.61s	17.15%	33,973.50s	17.03%	1,960.26s
90	10	4,929.50s	16.19%	35,324.40s	16.11%	6,670.96s
90	12	n/a	n/a	n/a	n/a	n/a
100	5	468.03s	13.46%	6,736.09s	13.31%	596.70s
100	8	4,474.56s	16.63%	39,533.50s	16.52%	2,587.09s
100	10	4,174.16s	15.24%	47,062.60s	15.19%	8,783.42s
100	12	n/a	n/a	n/a	n/a	n/a

Table 13 Deterministic vs. randomized solution of the constrained MC problem under general substitution sets and a cardinality constraint $|S| \leq C = 5$. All percentages are reported relative to the worst-case expected revenues of the optimal deterministic robust assortment.